


The Shape of the Disease

Alzheimer's on two axes — where connections are lost, and when

An introduction to the model being assembled from the published literature at
AdultCognitiveDisease.com



Benjamin Aaron Gustafsson

AdultCognitiveDisease.com

October 2026

“Alzheimer's is not an event. It is a process in time — a fifty-year sequence that begins decades before any symptom, and a fixed route through the brain in which one region is reached before the next.”



Alzheimer's is not an event

It is a process in time, and it occupies the better part of a human life. It begins in the third decade, in one nucleus of the brainstem no larger than a grain of rice. It ends in the eighth decade, in the inhibitory scaffolding of the cortex — fifty years and twelve centimetres away. Almost all of that is silent. By the time a person is sitting in a clinic being asked to remember three words, the process being named that day has been running for something close to forty years.

That claim is not new in its parts. Nearly every piece has been published, most of it decades ago, in literatures that do not cite one another. What is new is that it has been assembled. Adult Cognitive Disease is an open, privately funded reading of the published record — a century of it — carried out by frontier language models under a method we call Organic Network Synthesis, and returned for correction to the scientists whose programmes it interprets. Dozens whose work is cited have read their own sections, and the corrections have come at the margins rather than the core. Every identifier the site publishes — 3,480 DOIs and PubMed IDs — has been resolved against PubMed or Crossref and checked against the reference beside it.

What came out of that reading is a model with a definite shape, and this paper is an introduction to it. It has two axes, **time** and **space**: a fifty-year sequence beginning decades before any symptom, and a fixed route through the brain in which one region is reached before the next. It is a map rather than a verdict — the phases and the transitions between them are hypotheses, argued and graded where they appear — but it is a map of a kind the field has not had, and specific enough to be wrong in useful ways.

It is built from first principles on a single measurement: **where connections are lost, and when**. Not plaque counts, not tangle counts, not any one protein. The synapse is the unit, because synapse loss has been the strongest structural correlate of cognitive decline since Terry and colleagues established it in 1991 — and because findings from fields with no shared vocabulary, from brainstem catecholamine physiology to matrix biochemistry, can all be placed on the same two axes once each is read as a connection lost somewhere at some age. Disputes that look like disagreements about mechanism often turn out, once placed, to be disagreements about *timing*.

One note governs everything below: the orderings and ages that follow are **order claims, not measurements**. Nobody has watched a human brain lose its connections for sixty years. The sequence is the claim; the heights are illustrative.

The long silence


The first thing lost is not a memory. It is a connection, and the first connections to go belong to the **locus coeruleus** — roughly fifty thousand noradrenergic neurons in the dorsal pons that supply the entire cortex with the chemistry of waking and attention.

The neuropathology here is the most load-bearing import in the model and is not seriously contested. Braak and Del Tredici's staging places the earliest pretangle tau not in entorhinal cortex but in the coeruleus, and not in old age but in early adulthood: of 42 brains from individuals aged 4 to 29, 38 already carried it, and in 19 of 22 the material was confined to the coeruleus–subcoeruleus complex (Braak and Del Tredici, 2011).

Why there? The answer the model gives is bioenergetic. The coerulean neuron is among the most metabolically extravagant cells in the brain: autonomously pacemaking, firing throughout waking life, sustaining a vast thin branching arbor that reaches nearly the whole forebrain, and running a catecholamine chemistry that is itself a source of oxidative load. Such a cell lives permanently near the ceiling of its mitochondrial capacity, with no reserve to spend on slow unrepaired damage — so it is where the universal age-dependent erosion of cellular quality control first crosses from compensated to decompensated.

On this account the earliest-affected connections are approaching half gone by the time a person turns fifty. Nobody notices, because what the coeruleus does is not remembered — it is slept, attended to, and felt.

Figure 1 · The long silence




Connections from the locus coeruleus — the earliest place affected — plotted against age. The shaded stretch is the part of the disease nobody sees. On this account the earliest-affected connections are approaching half gone by the time a person turns fifty, and the diagnosis arrives near the far end of the curve rather than at its beginning.

A fixed route

The disease does not appear everywhere at once and it does not appear at random. It starts in one place and moves along the wiring to the next, in an order that is much the same from brain to brain — and that regularity is why it is possible to speak of a shape at all.


Figure 2 · A fixed route through the brain



The age at which each place has lost about a quarter of its connections. Read top to bottom, this is the order; read left to right, it is the timetable, and roughly twenty-two years separate the first station from the last. The rows are sorted by the curves themselves rather than by assertion — a place counts as reached when its own resident cells fail, not when the wiring arriving from the brainstem does.

The most consequential feature of that order is what sits at the top. The first two stations are not memory structures: they are small, ancient, modulatory nuclei, each sending a single thin axon an enormous distance to keep the rest of the brain in a working state. Both are in trouble before the structures they serve. The places that produce the recognisable symptoms are reached third, fourth and fifth, and the thinking surface last of all.

Put the six side by side and the disease stops looking like a switch and starts looking like a wave. Each keeps its own clock; the brainstem is well down its slope while the neocortex has barely started. There is no single moment at which someone gets Alzheimer's disease, and no single place where it lives. There is a front, and it moves.


Figure 3 · The wave, place by place

The same measure and the same axes, six times over. Grey lines are the other five places, repeated in every panel so that each can be read against the whole set; the coloured line is the panel's own. The curve moves rightward down the sequence: the brainstem is well down its slope while the neocortex has barely started.

And inside each place, it is selective

Zoom into one region and a further dimension appears. The disease does not thin out a piece of cortex evenly, like paint fading. It takes specific connections from specific kinds of cell and leaves others comparatively alone — for a while (Leng et al., 2021).

The clearest case is the **parvalbumin interneuron**, whose job is timing: it inhibits its neighbours in fast, precise rhythm and keeps cortical activity from running away with itself. It is wrapped in a protective lattice called a perineuronal net, and for six decades it is the best-protected cell of the five. Then the lattice is digested and it becomes the worst-affected one — losing, in the model's accounting, some 84 per cent of its connections against 54 per cent for the excitatory pyramidal cells around it.


Figure 4 · Not every cell falls together

The neocortex, broken out by the kind of cell losing its connections. For most of a life the parvalbumin interneuron loses less than anything around it. The crossing point, late on, is the moment its protection fails — and it is the argument for why the last phase of the disease arrives so much faster than the first.

That crossing point is the single most important event in the late disease. It is the moment the protection fails, and it is the model's explanation for why the last phase arrives so much faster than the first. (One correction belongs here, and the site carries it: a blind re-reading of the evidence found parvalbumin involvement to be late and functional, with **somatostatin** interneurons preceding it. The ordering claim survives; the identity of the terminal cell is held more loosely than it once was.)

Three phases, and the two bridges between them

That sequence is the phenomenology. Underneath them the model proposes a mechanistic spine: three phases, each load-bearing on a different substrate, a different cell type and a different decade — joined by two transitions that are specified rather than assumed.

Figure 5 · The arc: three phases and the two bridges between them

Each phase is load-bearing on a different substrate, a different cell type and a different decade; the gold joins are the transitions, which are specified mechanisms rather than gaps to be waved across. Decade boundaries are nominal — the phases overlap in any individual, and the ordering rather than the dating is the claim.

	Decade · place	The load-bearing failure
Phase I — Bioenergetic Ignition	3rd–5th · locus coeruleus and brainstem aminergic nuclei	Metabolic and custodial: PARP-1 hyperactivation, NAD ⁺ depletion, obstructed mitochondrial import, failing mitophagy
The Locus Coeruleus Bridge	the ascending LC projection	Two arms, one axonal arbor: withdrawal of the noradrenergic brake on microglia, and trans-synaptic tau seeding
Phase II — The Microglial Bridgehead	6th–7th · hippocampus	Collapse of the TGF- β /SMAD-maintained homeostatic microglial state; protective withdrawal and destructive attack become one act
The Proteolytic Turn	the aggrecan–brevican sheath	Three converging arms: a lipid-driven proteolytic switch (NLRP3 to IL-1 β to MMP-9), iron from ferroptotic oligodendrocytes, complement priming
Phase III — Synaptic Disintegration	8th+ · cortical and hippocampal interneurons and their nets	Structural and electrical: matrix digestion, loss of inhibitory competence, excitatory–inhibitory collapse

The bridges are the part we would most like other people to attack, because they are where a temporal theory earns the name. The first has a particular elegance: the very anatomy that makes the coeruleus a master regulator of the forebrain is what makes it the forebrain's instrument of harm. Noradrenaline is a standing brake on microglial activation, released diffusely across wide territories; the same ascending axons also carry seed-competent tau. When the nucleus fails, the hippocampus receives — through one arbor, in one epoch — both the cessation of the brake on its microglia and a stream of templated seeds. The loss of regulation and the delivery of pathology are one event. The second bridge then converts a cytokine-secreting microglial population into a matrix-digesting one, and removes the protection around the parvalbumin cell.


One further property is worth stating plainly: the architecture is **substrate-shifting**. The primary lesion is metabolic in Phase I, immunological in Phase II, structural in Phase III. That is why no single-substance theory has ever fitted the whole disease — and why *which molecule is primary* may be the wrong question, since the answer changes with the decade.

The symptoms nobody counts

The three places reached first are not memory structures, and what their failure produces is not forgetfulness. It is broken sleep, apathy, and getting lost — the three complaints most reliably written off as ordinary ageing. They are not soft signs. In a cognitively normal person apathy carries roughly three times the odds of progressing to mild cognitive impairment and twice the odds of progressing to dementia (Fan et al., 2021); in someone who already has mild cognitive impairment it raises the hazard of conversion by half again (Fresnais et al., 2023).


Broken sleep — the locus coeruleus. The nucleus that fails first sets the depth of wakefulness and the architecture of sleep. The direction of the claim matters and the model is careful about it: it does *not* argue that poor sleep causes the disease. The coerulean tau precedes any plausible sleep lesion. The lesion comes first, and disturbed sleep is among the earliest things it produces — after which it becomes a lever on what follows, because how badly sleep is disturbed forecasts the rate at which amyloid accumulates over subsequent years (Winer et al., 2020).

Apathy — the ventral tegmental area. A second small midbrain nucleus, tied to dementia in 1988 (Torack and Morris) and then left largely to Parkinson’s researchers. In amyloid-bearing mice its dopamine neurons begin dying at three months, before a single plaque forms — and only the mesolimbic ones; the neighbouring substantia nigra, which fails first in Parkinson’s disease, is untouched (Nobili et al., 2017). Every standard explanation of dopaminergic fragility predicts the opposite ranking. Dopamine sets the price an organism will pay for a goal, which is why its withdrawal reads as apathy rather than sadness, and it has a second consequence in the hippocampus, where it supplies something closer to a permission: a memory that has been formed is not yet a memory that will be kept, and whether it persists depends on a dopaminergic signal arriving in the hours around the event (Takeuchi et al., 2016). Encoding is not the step that fails. Both traces are laid down; only one is granted the signal that makes it survive the day. The conversation happened. By evening it is gone.

Figure 6 · The permission that is never given

The strength of a new memory trace across the day that follows it. Encoding is not the step that fails: both traces are laid down, and only one is granted the dopaminergic signal that converts it into something which survives the day. It is the shape of the complaint families actually report — the conversation happened, and by evening it is gone.

Getting lost — the somatostatin interneuron. The oldest finding in this disease that is not about acetylcholine is that somatostatin is depleted from the Alzheimer cortex (Davies, Katzman and Terry, 1980). It has never been overturned, and it has never been placed. The intuitive account — that these cells hold the brain's spatial metric — is wrong; silencing them leaves grid firing intact (Miao et al., 2017). What the cell does is more specific. A position computed from your own motion accumulates error, and that error is cleared on contact with the world: a wall, a doorway, a landmark (Hardcastle et al., 2015). A somatostatin interneuron sets the weighting between the input carrying the world and the input carrying memory (Leão et al., 2012) — the device that alternates between correcting by the world and computing from it. Health is not an accurate map; it is a repairable one. Lose the reset and the map is not erased, it becomes unresettable. The measurements agree: more than five thousand unassisted trips by nursing-home residents were 87 per cent *direct* travel, efficiency rather than quantity tracking cognitive status (Martino-Saltzman et al., 1991), and 325 community missing-person cases were “temporally appropriate but spatially disordered” (Rowe et al., 2011). The walking has an object. It is a search for the landmark that would reset the map.

Figure 7 · The map that cannot be reset

Error in the internal estimate of position, with and without the correction that clears it. Both lines drift at the same rate; the whole difference is the reset. Lose it and the map is not erased but unresettable — the landmark is seen, and cannot be bound to the position it ought to correct.

Where the independent lines meet

Laid on one set of axes, programmes that have never corresponded arrive at the same small number of places. The site names seven such convergence nodes; the densest is the **endosomal nexus**, where retromer failure, lipid peroxidation and autophagic collapse all land in one compartment. The plaque, on this reading, is not the agent but the residue of a process — a lesion with a developmental sequence, which is how Oskar Fischer first described it. His 1910 monograph set out eight morphological stages of the deposit and reported that the injury to the surrounding tissue tracked the stage rather than the number; two modern programmes re-derived that reading in molecular terms a century later.

One result of assembling them was not anticipated. Nine long-running programmes on the failing synapse each describe a cycle with a forward leg that makes, delivers or engages, and a return leg that recovers, releases, clears or disposes. In eight of the nine the reported lesion is on the **return** leg, and at six of thirteen stations the forward leg is measurably *up-regulated*, because the cell is compensating for a return that never arrives. Sorted that way the clinical record acquires a pattern: every intervention that suppressed a forward leg was null or harmful, and no intervention that restores a return leg has ever been tested in a human.

What the shape implies

If this account is even approximately right, most of what makes a person recognisably ill happens at the far end of this sequence — and almost every drug ever trialled has been given there, in the last few years of a forty-year process. The shape says the target is further left, and that interventions should be matched to the phase a patient is in rather than to the molecule a theory prefers. It also predicts, rather than excuses, the failure of single-target monotherapy: a substrate-shifting disease treated as though its target were fixed will disappoint at every phase but one.

It carries a hopeful implication too, and not a small one. What is being counted throughout is **connections**, not cells. A silenced neuron that has lost its inputs is in a different condition from a dead one, and the model argues that much of what is lost in the late phase is the former. Whether that is recoverable is an open question, and one of the most consequential in the field.

Two things should be said about how firmly this is held. Each claim is graded where it is argued, and the grades are not uniform: the brainstem neuropathology is strong, the bridges are hypotheses with named molecular arms, and two of the three phases have no instrument that can see them in a living person. And the scope is narrower than the word “Alzheimer’s” suggests — this is the disease biologically defined, which is a minority of dementia as it presents in clinic. The site says so, and maps the routes it does not describe.

An invitation

The reason to publish a model at this stage rather than when it is finished is that the people who can break it are reading it. If your work appears here it has a page, and that page states your programme at full strength before it tests it. Tell us where it is wrong, out of date, or graded too generously. Corrections are published with the objection kept under your name; dissent is the most valuable thing the project receives, and it is not edited into agreement.

Everything is open and free — 827 interlinked articles, 635 concepts, the seven monographs, the papers in full, and the charts above in interactive form, at **AdultCognitiveDisease.com**. The place to start is *The Temporal Architecture of Collapse*, the paper the rest of the model elaborates.