THE HOMEWARD VECTOR

Why the Person With Dementia Walks — Wandering as a Correctly Executed Search Program Running on a Map That Can No Longer Be Re-Anchored, and Why the Somatostatin Governor's Failure Is a Failure of the Reset Rather Than of the Metric

The Metric and the Reset · The Switch, Not the Signal · The Program the Ant Runs
Wandering Is Not Random · Why a Receptor Cannot Supply an Address

A Dissertation Submitted in Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy in Neuroscience

Prepared under the Organic Network Synthesis methodology

AdultCognitiveDisease.com

Benjamin Aaron Gustafsson

August 2026

The behavioural companion to The Unnetted Governor, which asked what the somatostatin neuron does inside the circuit and declined the grid-cell connection it was offered. This volume asks what the same lesion does out in the corridor, and finds the connection real one level down — at the operation that keeps any code usable: the reset.

"No animal carries an accurate map. Every reckoning drifts, in every nervous system that has one, and it stays usable only because the world is allowed, several times a second, to correct it. Health is not a good map. Health is a map that can be fixed."

— ONS Methodology Working Notes, 2026

For everyone who has been told that the walking means nothing — and for those found in the woods, less than a mile from where they started, which is not the geometry of someone who wandered off, but the geometry of someone who searched in the wrong place.

THE COLLAPSE FRAMEWORK THE INHIBITORY LINE

The cell — *The Unnetted Governor*

The kinase it cannot restrain — *The Coerulean Pincer* · *The Silenced Eraser*

The lattice it does not wear — *The Architect's Scaffold* · *The Coerulean Shears*

The nucleus that shares its exposure — *The First Ember* · *The Coerulean Clock*

The behaviour ← The Homeward Vector — this volume

The corpus had written the cell, the kinase, the lattice and the nucleus. It had not written what any of them look like from the corridor. This volume supplies the behavioural reading, and finds that the reading requires refusing the connection the question arrived with. The somatostatin neuron does not sharpen the grid — a clean pharmacogenetic dissociation says the hexagonal code belongs to its parvalbumin sibling. What it holds instead is the switch that admits the world into the running map, in the one lamina where the world arrives; and the loss of that switch does not erase the map, it makes the map unresettable. Wandering is then not the absence of a program but the running of one — the systematic search any path-integrating animal executes when its home vector expires — in a system that can no longer register that it has arrived.

Contents

Abstract

A person with dementia walks. She has been walking for forty minutes, along a corridor she has lived in for three years, past a door with her own name on it, and she will walk past it again. The clinical vocabulary for this is *wandering*, a word that carries the assumption that the walking is without object. This dissertation argues that the assumption is false, that the falsity is measurable, and that the mechanism which makes the walking purposeful is a specific and identifiable failure in the entorhinal–hippocampal circuit — one that the corpus has already described from the inside, in its somatostatin volume, without noticing what it looks like from the outside.

The dissertation begins by stating the popular account it was asked to develop, because that account is nearly right and its errors are instructive. In the popular account, somatostatin interneurons die early; their death disinhibits grid cells; the hexagonal metric degrades into noise; path integration fails; the internal map is lost; anxiety follows; and the person walks in search of a landmark that will let the map be re-set. Every clause of that chain is either true or nearly true, except the second and third — and those are the load-bearing ones. This volume's companion, *The Unnetted Governor*, established from a clean pharmacogenetic dissection that silencing somatostatin interneurons of the medial entorhinal cortex has *no effect* on grid cells, while silencing parvalbumin interneurons antagonises the hexagonal code (Miao and colleagues, 2017). The anatomy agrees: entorhinal somatostatin cells preferentially inhibit layer III–V pyramidal cells and largely spare layer II, where grid cells live (Kecskés and colleagues, 2020). The grid's inhibitory partner is not this cell. The popular chain breaks at its second link.

The dissertation's first move is to show that the chain's *conclusion* nonetheless survives, by a different and better route, and that the route runs through a fact the popular account omits entirely. Grid cells do not merely encode position; they accumulate error while doing so, and that error is corrected by contact with the environment. Grid firing patterns drift coherently and are *reset* by encounters with environmental boundaries (Hardcastle and colleagues, 2015), and removal of visual landmarks profoundly impairs grid periodicity (Pérez-Escobar and colleagues, 2016). Path integration is therefore not a self-sufficient computation. It is a computation with a scheduled correction, and the correction is the part that requires the world. Health is not an accurate map. Health is a *repairable* map.

The dissertation's second and signature move is to locate the repair operation anatomically and to show that it is the somatostatin neuron's, even though the metric is not. The somatostatin-expressing oriens–lacunosum-moleculare interneuron of CA1 differentially weights the two great inputs to the pyramidal cell: it reduces the influence of the extrahippocampal input from entorhinal cortex while facilitating the intrahippocampal input from CA3 (Leão and colleagues, 2012). That is, in circuit terms, a switch between *correction by the world* and *computation from memory* — the alternation the theta-phase models of encoding and

retrieval require (Hasselmo and colleagues, 2002). The reset is not a signal. It is a scheduled alternation, and the somatostatin cell is the device that schedules it, in the one lamina where the entorhinal signal actually arrives. Lose the governor and the alternation collapses: the external channel is no longer weighted against the internal one, and a landmark that is seen cannot be bound to a map that is being computed. The map is not erased. It becomes unresettable.

The third move takes the account outdoors and tests it against what is actually measured about wandering, and here the evidence is stronger than the field's own vocabulary suggests. Video observation of more than five thousand unassisted travel events in nursing-home residents found travel to be 86.8 per cent *direct*, 11.6 per cent lapping, 0.7 per cent pacing, and 0.9 per cent random — and travel efficiency, not travel quantity, tracked cognitive status (Martino-Saltzman and colleagues, 1991). Wandering is overwhelmingly goal-directed locomotion whose goal is not being reached. Analysis of 325 community missing incidents found them "temporally appropriate but spatially disordered," typically beginning as an ordinary permitted errand rather than as a wandering episode, and lacking the repetitive signature of wandering altogether (Rowe and colleagues, 2011). Two behaviours, then, and this dissertation insists on separating them: *wandering*, the repetitive search inside a familiar enclosure, and *getting lost*, the single-trip navigational failure outside it. One lesion, two environments, two expressions. The corpus should not fuse them, and the clinical literature's own data say so.

The fourth move supplies the comparative anchor, in the same spirit in which the companion volume used experimental epilepsy to establish that the somatostatin cell's vulnerability is a property of the cell type rather than of amyloid. When a desert ant reaches the end of its home vector and does not find the nest, it does not stop and it does not move at random: it executes a stereotyped systematic search of loops of ever-increasing size, centred on the point where the nest should have been (Wehner and Srinivasan, 1981). Search after path-integration failure is a *program*, not a breakdown — an ancient, structured, centre-biased routine that runs when the reckoning runs out. This dissertation's proposal is that human wandering is that program, released correctly, in a system that can no longer register the program's success condition. The search does not terminate because the operation that would terminate it — binding a recognised landmark to the running map — is the operation that has failed.

Around this spine the dissertation assembles what the account requires: the two clocks of failure (the somatostatin lesion is preclinical, while sixty per cent of entorhinal layer II neurons are already gone at a Clinical Dementia Rating of 0.5, so the reset fails *before* and the metric fails *after*); the noradrenergic alarm that converts a navigational error into an urgent one; the circadian phase delay that gives the search an hour; and the cholinergic arm, which is a single denervation with two navigational consequences, because muscarinic blockade reduces grid tuning while the loss of septo-hippocampal drive is the identified mechanism of somatostatin-cell dysfunction.

The volume then makes two corrections to the diagnostic and therapeutic frontier it was asked to endorse. On virtual reality: immersive path-integration testing genuinely outperforms

standard cognition, discriminating biomarker-positive from biomarker-negative mild cognitive impairment with an area under the curve of 0.90 against 0.57 for the best cognitive test (Howett and colleagues, 2019), and the deficit is selective and present decades early in risk carriers (Bierbrauer and colleagues, 2020). But the same laboratory that built the case found that virtual navigation measures could *not* reliably identify which patients would become spatially disoriented in their own community (Puthusseryppady and colleagues, 2022b). The tests detect the disease; they do not yet detect the wanderer. On pharmacology: the proposal to restore inhibitory tone with a subtype-selective GABA-A modulator founders on an address problem that the receptor literature has already resolved and the navigation literature has not read. The alpha-5 subunit is preferentially targeted not to pyramidal dendrites but to the synapses made by vasoactive-intestinal-peptide and calretinin terminals *onto somatostatin interneurons themselves* (Magnin and colleagues, 2019) — which are disinhibitory inputs. A positive modulator at that address does not restore the governor; it amplifies the input that silences it. The pharmacology of tone cannot supply the anatomy of address.

Every joint is graded in a three-tier validity ledger. Six falsifiable predictions are stated, including the one that would kill the thesis outright. And the care corollary is drawn plainly, because it is the practical yield of the whole argument: if wandering is a search program that cannot register success, then restraining the search treats the symptom and starves the mechanism, and the humane and mechanistically correct intervention is to supply, in the built environment, the unambiguous terminating landmark that the broken reset can still, sometimes, accept.

I. The Question, and the Chain It Assumes

The question this dissertation was set is a good one, and it arrives with an answer already attached. It is worth writing that answer out in full, because it is the most widely circulated mechanistic account of wandering in dementia, because it is largely correct, and because the two places where it fails are precisely the places where the interesting biology is.

The received chain. In the received account the argument runs: somatostatin-expressing inhibitory interneurons are among the earliest cell populations lost in Alzheimer's disease; their loss disinhibits the grid cells of the medial entorhinal cortex; the hexagonal firing lattice degrades into noise; path integration — the running vector sum of distance and direction that lets an animal know where it is without looking — fails; the internal coordinate system collapses; even a familiar bedroom becomes unlocatable; the resulting spatial amnesia generates intense anxiety; and the anxiety releases an evolutionarily ancient drive to move, to find a salient landmark, and to re-anchor the map. Wandering, on this account, is not confusion. It is a rescue operation.

What is right about it, stated first. Three of the chain's claims are secure and this dissertation adopts them without qualification. Somatostatin interneurons *are* lost early: the largest available human multimodal atlas resolves the disease into an early phase — slowly increasing pathology, inflammatory microglia, reactive astrocytes, loss of somatostatin-positive inhibitory neurons — and a later phase in which excitatory neurons and the parvalbumin and vasoactive-intestinal-peptide interneuron classes fall, with the early-phase donors cognitively unimpaired (Gabitto and colleagues, 2024). Grid-cell failure in Alzheimer models and in human risk carriers *is* real, early, and well documented, at the earliest fibrillar amyloid deposition, in tauopathy, and decades before disease in young carriers of the apolipoprotein E ε4 allele (Ying and colleagues, 2022; Fu and colleagues, 2017; Ridler and colleagues, 2020; Kunz and colleagues, 2015). And path-integration failure *is* one of the earliest measurable cognitive signs of the disease, more discriminating than the standard neuropsychological battery (Howett and colleagues, 2019; Bierbrauer and colleagues, 2020). The chain's premises and its conclusion are in good order.

Where it fails, and why that matters. The failure is at the joint. The claim that somatostatin-cell loss disinhibits grid cells and thereby blurs the hexagonal lattice is not supported, and the experiment that tests it directly returns the opposite result. In freely moving mice, pharmacogenetic silencing of parvalbumin interneurons antagonised the hexagonal spatial selectivity of grid cells, most strongly in layer II, and reduced speed modulation in co-localised speed cells; silencing somatostatin interneurons had no impact on grid cells or speed cells at all, and instead decreased the spatial selectivity of cells with discrete, aperiodic firing fields (Miao and colleagues, 2017). The anatomy predicts this: somatostatin interneurons of the medial entorhinal cortex preferentially inhibit layer III–V pyramidal cells and leave layer II largely unaffected (Kecskés and colleagues, 2020), and layer II is where grid cells are densest. The companion volume *The Unnetted Governor* stated this correction and declined the connection. This volume inherits the correction and does not re-litigate it.

The question this dissertation asks instead. If the somatostatin cell does not sharpen the grid, and yet the somatostatin lesion is the earliest of the disease's inhibitory lesions, and yet the behaviour that most defines living with dementia is a navigational behaviour — then either the somatostatin lesion is irrelevant to wandering, or the chain is routed through the wrong navigational function. This dissertation argues the second, and its argument turns on a property of the spatial system that the received account does not mention at all: the map has to be *corrected*, continuously, by the world, and the correction is a distinct operation from the computation it corrects. Grid cells accumulate error. Something resets them. The reset is not the metric, and the reset is where the governor lives.

The shape of the answer, and the shape of the volume. The dissertation will argue that wandering in dementia is the behavioural expression of a failed *reset*, not of a failed *metric*; that the reset in CA1 is implemented by the alternation between entorhinal and intrahippocampal input

which the somatostatin oriens–lacunosum-moleculare cell demonstrably controls; that this alternation is degraded very early, in the preclinical phase in which the somatostatin cell is lost, and before the layer II grid population is destroyed; that the resulting behaviour is not the absence of a program but the running of one — the systematic search that any path-integrating animal executes when its home vector expires; and that the search does not stop because the system can no longer register that it has arrived. Sections II and III establish the metric and the reset. Sections IV to VII locate the lesion. Sections VIII to XI take the argument outdoors and test it against the wandering literature. Sections XII to XIV supply the alarm, the clock, and the cholinergic arm. Sections XV and XVI make two corrections to the diagnostic and therapeutic frontier. The remainder grades, predicts, and draws the care corollary.

II. The Metric What the Grid Computes, and What It Costs

Before the failure can be described, the machine must be. This section is deliberately brief on the parts that are famous and deliberate on the parts that are not.

The lattice. A grid cell of the medial entorhinal cortex fires at multiple locations in an environment, and those locations form a regular triangular lattice tiling the available space (Hafting and colleagues, 2005). This is a different object from a place cell, which fires at one location: a place cell says *here*, a grid cell supplies a coordinate system in which *here* can be expressed. Cells of a given grid module share spacing and orientation and differ in phase; modules of different scale are stacked; and the combination gives a metric — a means of expressing displacement, not merely position. The discovery is thirty years downstream of the place cell and it earned a Nobel Prize because it supplied the missing half of the cognitive map: not the label, but the ruler.

Path integration, and its price. The reason a metric matters is that it permits dead reckoning: an animal that can integrate its own velocity over time can know its displacement from a starting point without seeing anything at all, and can therefore compute the vector home. The grid network is the best candidate substrate for that integration. Selective disruption of grid-cell firing, achieved by removing NMDA-type glutamate receptors from the retrohippocampal region without affecting other spatially selective cells in medial entorhinal cortex or hippocampus, impairs path-integration performance (Gil and colleagues, 2018). This is the causal link the behavioural literature needs, and it is the strongest one available: break the grid, and dead reckoning breaks with it.

The grid's recurrent connectivity is inhibitory, and this is not a detail. Layer II stellate cells of the medial entorhinal cortex — the principal grid population — are almost entirely unconnected to one another by direct excitatory synapses. Their recurrent interaction is disynaptic, through fast-spiking inhibitory interneurons, and a network of stellate cells coupled by such inhibition is sufficient to generate grid-like firing (Couey and colleagues, 2013). The grid is, in a precise structural sense, an inhibitory network. That fact is what makes the popular chain's instinct correct in general and wrong in particular: the hexagonal code really does depend on interneurons for its existence, not merely for its polish — but the interneurons in question are the fast-spiking, perisomatic, parvalbumin-class cells whose silencing degrades the code (Miao and colleagues, 2017), not the dendrite-targeting somatostatin cells whose silencing does not.

The self-motion inputs, and their corruption. A path integrator needs speed and direction. Both are represented in the medial entorhinal cortex, and both are vulnerable. In the rTg4510 tauopathy model, grid firing patterns are largely absent and the neural representation of running speed is significantly disturbed, with the normally tight coupling between locomotion and theta- and gamma-band local field potential power abolished (Ridler and colleagues, 2020). Speed coding is not an accessory to the grid; it is the integrand. A metric whose velocity input is corrupted does not merely become imprecise — it becomes *systematically* wrong, and the error compounds with distance travelled. This is worth holding in mind for Section IX, because the ant that overshoots its nest and the person who walks past her own door are, formally, the same kind of failure.

The human case, and the honest caveat. Grid-like activity exists in humans. It has been recorded directly from medial temporal lobe neurons during virtual navigation (Jacobs and colleagues, 2013), and it is detectable non-invasively as a hexadirectional modulation of the entorhinal blood-oxygen-level-dependent signal, higher for movement aligned with grid orientation than for movement misaligned with it (Doeller and colleagues, 2010). This dissertation will lean on the human grid literature in Section XV, and it therefore states the caveat now: the hexadirectional functional-imaging signal is an aggregate proxy, not a recording of grid cells, and a reduction in that signal is compatible with several underlying changes. It is the best non-invasive index available and it should not be described as a measurement of grid-cell firing.

III. The Reset — The Operation the Received Account Omits

This is the pivotal section of the dissertation, and its content is a single fact with a large consequence.

Path integration is lossy, and the loss is continuous. An integrator accumulates error. Every biological velocity estimate is noisy; the noise integrates; and the position estimate drifts. This is not a pathological property, it is a mathematical one, and it applies to a healthy young grid network exactly as it applies to a diseased one. The consequence is that no purely internal spatial system can remain accurate for long, and therefore that any usable spatial system must include a means of correction from outside itself.

The correction is measured, and it happens at boundaries. The demonstration is direct. Grid-cell firing patterns accumulate error and drift coherently as an animal moves, and that error is *corrected* upon encounters with environmental boundaries — the drift is reset by contact with the walls of the enclosure, and grid error is, in the authors' framing, literally bounded by the environment (Hardcastle and colleagues, 2015). The grid is not an autonomous coordinate system that the world occasionally consults. It is a coordinate system that would degrade within minutes if the world did not repeatedly repair it.

The correction also depends on distal landmarks. Boundaries are not the only anchor. Removal of visual landmarks causes a profound impairment of grid-cell periodicity; the speed code of medial entorhinal neurons changes in darkness; and the activity of border cells becomes less confined to environmental boundaries (Pérez-Escobar and colleagues, 2016). Landmarks do not merely orient the grid; they sharpen the metric itself. And the anchoring is plastic: entorhinal grids rescale with experience when the geometry of a familiar environment is altered (Barry and colleagues, 2007), which is to say that the mapping between the internal lattice and the external world is a learned, updatable registration rather than a fixed one.

The claim, stated. Spatial competence is therefore *two* operations, not one. There is a metric, which computes displacement from self-motion and drifts; and there is a reset, which periodically re-registers that metric against a stable external reference and discards the accumulated error. Both are necessary. Neither is sufficient. A person with a perfect metric and no reset will be lost within minutes of the last correction; a person with a degraded metric and an intact reset will make constant small errors and constantly repair them, and will not get lost at all. This is the reason a healthy person can walk an unfamiliar city, integrating badly, and still arrive: not because the reckoning is good, but because the repair is frequent.

Why the omission matters for the received chain. The received account treats the loss of navigation as the loss of the map, and therefore predicts that the remedy is a landmark: give the person a strong salient cue and the map will re-anchor. That prediction is testable and, as Section VIII will show, it does not survive contact with the data. If the deficit were in the metric alone, the environment would rescue it — the person would drift, encounter a wall or a doorway, reset, and continue. That she does not is the central clinical fact, and it points at the reset rather than at the metric. A landmark is only useful to a system that can still perform the operation of binding it. The

received account has identified the right symptom and the wrong broken part.

A note on vocabulary. This dissertation will use "reset" in the specific sense established here — the discrete re-registration of an internally computed position estimate against an external reference — and not in the loose sense of "recovering one's bearings." The distinction matters because the loose sense is a psychological description and the specific sense is a circuit operation with an address, which Sections V and VI supply.

IV. The First Correction The Governor Does Not Sharpen the Grid

The companion volume established this and graded it, and the present volume must state it plainly before it can build on it, because the whole architecture of the argument depends on the correction being accepted rather than quietly evaded.

The experiment. Pharmacogenetic silencing during recording in freely moving mice dissociated the two great inhibitory classes of the medial entorhinal cortex. Silencing parvalbumin cells antagonised the hexagonal spatial selectivity of grid cells, especially in layer II, and reduced the speed modulation of co-localised speed cells. Silencing somatostatin cells had no impact on grid cells or speed cells; what it degraded was the spatial selectivity of cells with discrete, aperiodic firing fields. Border cells and head-direction cells were unaffected by either manipulation (Miao and colleagues, 2017). This is a clean positive-and-negative result within a single study, which is the strongest form such a dissociation can take.

The anatomy that predicts it. Somatostatin-positive interneurons of the medial entorhinal cortex preferentially inhibit layer III–V pyramidal cells and leave layer II neurons largely unaffected; their inhibition is dual, carried by gamma-aminobutyric acid and by the somatostatin peptide itself; it is prolonged rather than fast; and behavioural work assigns the motif to the modulation of working-memory formation rather than to the retrieval of learned routes (Kecskés and colleagues, 2020). The cell is a deep-layer controller. It does not have its hands on the grid layer. Two studies from different laboratories, one anatomical and one physiological, published three years apart, agree.

What the correction costs, and what it buys. It costs the tidiest available story. It buys three things. First, accuracy, which is not a small return in a literature where the somatostatin finding has been repeatedly over-read since 1980. Second, a sharper target: if the hexagonal code is

the parvalbumin cell's and the parvalbumin cell falls late (Gabitto and colleagues, 2024), then acute grid degradation is a *late* lesion in the inhibitory sequence, and the earliest inhibitory lesion must express itself through some other spatial function. Third — and this is the productive consequence — it forces the question of what the somatostatin cell *does* hold, and the answer, once looked for, turns out to be more relevant to wandering than the grid is.

What the correction does not license. It does not license the conclusion that the somatostatin cell is irrelevant to spatial failure in Alzheimer's disease. Three routes remain open, and the companion volume named them: the cell governs a different space-coding network, the aperiodic one; it governs the deep-layer output stage, which is what the entorhinal cortex *sends* rather than what it computes; and its early loss degrades the chronic excitatory–inhibitory conditions under which the layer II grid population must survive. The present volume adds a fourth, which it regards as the most consequential, and which requires leaving the entorhinal cortex for the hippocampus.

V. The Switch, Not the Signal What the Governor Actually Holds

Here the dissertation makes its signature claim. It is assembled from measured parts that belong to two literatures — hippocampal circuit physiology and spatial-memory theory — and it is not, so far as the present search found, stated anywhere as such.

The first fact: the somatostatin cell weights the two inputs to CA1 against each other. The CA1 pyramidal cell receives two great excitatory inputs. One arrives from CA3 by the Schaffer collaterals onto proximal apical dendrites in stratum radiatum, and it carries intrahippocampal information — the associative, pattern-completing, recall-like signal. The other arrives from the entorhinal cortex by the perforant path onto distal apical dendrites in stratum lacunosum-moleculare, and it carries extrahippocampal information — the current sensory and spatial state of the world. Optogenetic dissection established that the somatostatin-expressing oriens–lacunosum-moleculare interneuron differentially weights these two channels: it facilitates transmission of the intrahippocampal CA3 input while reducing the influence of the extrahippocampal entorhinal input, and it is directly driven by subcortical cholinergic afferents (Leão and colleagues, 2012). The oriens–lacunosum-moleculare cell is, in one sentence, a switch on the world-input channel.

The second fact: the two channels are supposed to alternate. The theoretical framework that gives this switch a function is the theta-phase model of encoding and retrieval, in which the hippocampal formation alternates rapidly between a condition favouring encoding — strong entorhinal input to CA3 and CA1 — and a condition favouring retrieval — strong CA3 input to CA1 — with the alternation carried on the theta cycle, and with the prediction that the preferred theta phase of CA1 spiking should differ for information being encoded and information being retrieved (Hasselmo and colleagues, 2002). The system is not supposed to run on the world, and it is not supposed to run on memory. It is supposed to switch between them several times a second, so that what it computes internally can be checked against what is actually there.

The joined claim. The reset of Section III — the discrete re-registration of an internally computed position against an external reference — requires exactly this alternation. To correct a running estimate you must be able to hold the estimate (internal mode) and then admit the correction (external mode) and then resume holding the corrected estimate. A system stuck in internal mode cannot be corrected; a system stuck in external mode has nothing stable to correct, because each moment's sensory input simply overwrites the last. The operation is not a signal that arrives; it is a *scheduled alternation*, and the somatostatin oriens–lacunosum-moleculare cell is the device that schedules it in the lamina where the external signal lands. This dissertation's claim is therefore that the somatostatin governor's territory is not the metric but the *timing of the metric's correction*, and that its loss produces neither a blind map nor a blank one, but an unresettable one.

Why this is the right shape of lesion for the clinical picture. An unresettable map predicts two things that a merely-degraded map does not. It predicts that the environment will fail to rescue the person — that a familiar doorway, a photograph, a name plate, will be seen and will not help, because seeing is not the failed step and binding is. And it predicts *instability* rather than *absence*: the same room will not be reliably the same room across visits, because the registration between the internal lattice and the external scene is being re-made from scratch and badly each time rather than being maintained. Both are recognisable to anyone who has watched the disease. Neither follows from grid blur alone.

Grading the claim honestly, before building further on it. The Leão result is a measurement, obtained by optogenetic manipulation, of the direction in which oriens–lacunosum-moleculare activity shifts the balance of the two inputs. The Hasselmo framework is a model with substantial supporting evidence but it is a model. And no experiment cited here shows that removal of somatostatin cells specifically abolishes the *reset* operation as such, in a navigating animal, measured as failure to correct accumulated path-integration error at a boundary. That experiment has not been done, and Section XVIII states it as the load-bearing prediction of the volume. The ledger grades this section as a synthesis of measured parts with one unmeasured joint — the same grade, and for the same structural reason, that the companion volume assigned to its coincidence of address.

VI. The Lamina Where the Reset Would Happen, If It Happened

The companion volume's signature contribution was an observation about a single hippocampal layer. That observation was made in the service of an amyloid argument. It has a second reading, in the service of a navigational one, and stating it is the work of this short section.

Three functions in one lamina, and now a fourth. Stratum lacunosum-moleculare is where the entorhinal perforant path terminates on the distal dendrites of CA1 pyramidal cells; it is where the somatostatin oriens–lacunosum-moleculare interneuron sends its axon; and it is where the somatostatin-receptor-dependent neprilysin capacity that degrades amyloid-beta is localised, since deletion of the somatostatin receptor subtypes 1 and 4 reduces presynaptic neprilysin specifically in that layer (Nilsson and colleagues, 2026; Leão and colleagues, 2012; Lovett-Barron and colleagues, 2012). The companion volume named this coincidence and built its amyloid loop on it.

The navigational reading. If the reset is the admission of the external signal into the running internal estimate, then stratum lacunosum-moleculare is the physical site of the reset in CA1: it is the only place where the entorhinal signal arrives, and the gate on it is the governor's axon. This yields a claim of unusual economy, and one that is either a real feature of the architecture or a genuine coincidence. The lamina in which the disease's defining molecule is disposed of is the lamina in which the map is corrected, and the same cell operates both. An injury confined to that thin layer is simultaneously an injury to amyloid clearance and an injury to spatial re-anchoring. There is no need to postulate two lesions to explain a patient who has both rising amyloid and a failing sense of place.

The measured consequence of a lesion at this junction. It has been observed. In an amyloid-precursor-protein knock-in model, hippocampal CA1 remapping was disrupted while CA1 spatial responsiveness was relatively preserved, entorhinal spatial tuning was severely lost, and the fast gamma oscillations that couple the entorhinal cortex to CA1 were substantially impaired (Jun and colleagues, 2020). Read against Section V, that combination is diagnostic. *Remapping* — the reassignment of the hippocampal representation when the environment changes — is precisely a re-registration operation; *spatial responsiveness* is the representation itself. The model preserved the representation and broke the re-registration, and it broke the entorhinal-to-CA1 coupling that carries the external signal. That is the predicted signature of a failed reset with an intact map, and this dissertation regards it as the strongest existing indirect support for the claim of Section V.

What is still missing. The Jun result is a model result, and it does not manipulate the somatostatin cell. The inference from "entorhinal-to-CA1 coupling fails and remapping fails" to "the governor's gate on the lamina has failed" is an inference, not a measurement. It is graded as such.

VII. Two Clocks What Fails First, and What That Predicts

The corpus's standing method is to ask not only what breaks but in what order, because order carries therapeutic content. Two timetables run here, and they are decades apart.

The first clock: the reset. The somatostatin interneuron is lost in the early, cognitively silent phase of the human disease, before excitatory neurons and before the parvalbumin and vasoactive-intestinal-peptide interneuron classes, and its loss precedes tangle deposition in the same tissue (Gabitto and colleagues, 2024). Before the cell is lost, its function is already deranged: in an amyloid model, two-photon imaging of the oriens–lacunosum-moleculare interneuron found severely impaired synaptic rewiring at both input and output, with learning-dependent remodelling disrupted and reduced cholinergic drive identified as a critical mechanism (Schmid and colleagues, 2016). If Section V is right, the reset therefore begins to fail in a person who is, by every clinical instrument then available, normal.

The second clock: the metric. The grid population itself is destroyed on a different and later schedule — but not by much, and the human number is startling. Stereological counting in the entorhinal cortex of individuals with a Clinical Dementia Rating of 0.5 — very mild impairment — found thirty-two per cent fewer entorhinal neurons overall than controls, and *sixty per cent* fewer in layer II specifically; in severe disease the layer II loss approaches ninety per cent (Gómez-Isla and colleagues, 1996). Layer II is the grid layer. The metric's substrate is therefore being destroyed outright, in humans, at the very mildest clinically detectable stage — not blurred by disinhibition, but removed.

What the two clocks together imply. Three things follow, and they matter for the rest of the volume. First, the popular chain's causal direction is inverted at its centre: it is not that inhibitory failure blurs the grid and thereby produces spatial amnesia; it is that the grid layer is being killed on its own account, while the earlier and quieter lesion is to the operation that would have compensated for a degraded grid. Second, the compensation is exactly what a healthy person relies on — Section III's point that health is a repairable map, not an accurate one. A person losing layer II neurons but retaining an intact reset would be a person who navigates inefficiently and repairs constantly. A person who has lost the reset first and then begins to lose layer II has no such margin.

Third, the order predicts the clinical sequence, and the sequence is what is observed: subtle path-integration inaccuracy long before disorientation, then disorientation in unfamiliar places, then disorientation in familiar ones, then the corridor.

The uncomfortable corollary about which cell is "first." This dissertation does not claim that the somatostatin lesion causes the layer II loss. The evidence does not support that and the companion volume graded the analogous claim as inference. What it claims is narrower: the two lesions are not redundant, they are sequential, and the earlier one removes the tolerance that would otherwise have concealed the later one. That is a claim about *reserve*, and it is consistent with the corpus's standing position that resilience — not mechanism — is the anchor the human evidence most strongly supports.

VIII. The Evidence That the Broken Part Is the Reset

An argument of this shape must be exposed to disconfirmation from the clinical record rather than only supported from the animal record. Three lines of evidence bear on whether the deficit is in the map or in the map's repair, and all three point the same way.

The familiar-room problem. The received account's own strongest illustration is its weakest point. It observes, correctly, that a person with dementia may fail to recognise her own bedroom. But a metric failure alone cannot produce that. Recognising a room does not require path integration; it requires matching a scene to a stored representation, which is a different operation supported by different structures. What a metric failure produces is *displacement* error — arriving in the wrong place, overshooting, being unable to return. What the familiar-room phenomenon indicates is that the scene, even when it is available and even when it is recognised, is not being used to place the person within a spatial frame. That is a binding failure, not a mapping failure, and binding is the reset.

The landmark-density null result. The prediction that follows from a metric-only account is that landmark-rich environments should protect. It has been tested. Community-dwelling patients with Alzheimer's disease and age-matched controls were tracked by global positioning system over two weeks; any spatial disorientation during the tracking period was recorded; and a spatial-buffer methodology captured the outdoor landmark density and road-network features of the environments each participant actually visited. The environments visited by patients who experienced disorientation did not significantly differ, in outdoor landmark density or in road-network structure, from those visited by patients who remained oriented (Puthusseryppady and

colleagues, 2022a). Landmark abundance did not separate the disoriented from the oriented. This is precisely the null result a reset account predicts and a metric-plus-landmark account does not: the anchors were there, and they did not help, because the operation that consumes anchors is the broken one.

The behavioural adaptation the same study found. The same tracking study found that patients, when unaccompanied, sharply restricted their own spatial and temporal range — fewer outings, less time moving, shorter distances, smaller radius from home — while navigating like controls when accompanied (Puthusseryppady and colleagues, 2022a). This is not a mechanistic finding, but it is a revealing one, and it belongs in a volume about purposeful behaviour: patients are *managing* a known deficit by shrinking the domain over which correction is required. A person with an unresettable map and enough insight will contract her world to the radius over which drift stays tolerable. That is the same logic, exercised deliberately, that Section IX will describe an ant exercising reflexively.

A rival explanation the ledger will record. There is an alternative reading of the familiar-room phenomenon: that the landmark is not recognised at all, a deficit of object and scene identity rather than of spatial binding. The division of labour in the entorhinal cortex — with the medial division carrying spatial and self-motion information and the lateral division carrying object, item and contextual information — makes this entirely plausible, and the two deficits would look similar from the corridor (reviewed in Coughlan and colleagues, 2018). This dissertation does not claim the binding account excludes the identity account; it claims the binding account is required, because landmark identity alone cannot explain the landmark-density null result, and grades the joint accordingly.

IX. The Program Vector Expires — What an Animal Does When the Home

The companion volume anchored the somatostatin cell's vulnerability in a different disease, and gained from the independence of that anchor. This volume anchors the behaviour in a different phylum, and gains the same way.

The desert ant, and why it is the right comparison. *Cataglyphis* forages across featureless salt pan in temperatures that make error fatal, and returns to a nest entrance that is a hole in the ground with no local landmark. It navigates by path integration: it maintains a running home vector, and when it turns for home it runs that vector off. The species is the classic experimental

preparation for dead reckoning precisely because it has so little else to go on. It is therefore the cleanest available model of what a nervous system does when the reckoning fails.

What it does is a program, not a collapse. If a homing ant reaches the end of its home vector and the nest is not there, it does not stop, and it does not move at random. It executes a stereotyped systematic search: loops of ever-increasing size, beginning and ending at the point where the nest ought to have been, and pointing in successively different azimuthal directions, so that the centre — where the nest is most probable — is searched most intensively and the periphery least (Wehner and Srinivasan, 1981). The structure is not incidental. It is an approximately optimal solution to searching under a known error distribution, and it is executed identically by naive animals.

The three properties of the program, and their behavioural signatures. The systematic search has three features worth naming, because each has a counterpart in the human clinical description. It is *centre-biased*: the animal keeps returning to the estimated goal rather than proceeding outward monotonically — the behavioural signature of which is repetition and looping. It is *persistent*: it continues far past the point at which the vector has been exhausted, because stopping is worse than searching. And it is *terminated by success*, not by time: the program ends when the goal is found, and it has no other stopping rule.

The transposition, stated as a hypothesis. This dissertation proposes that human wandering in dementia is this program — an ancient, structured, centre-biased search released by the expiry of a positional estimate — running in a nervous system that has lost the operation by which the program's success condition is evaluated. The person is not searching aimlessly; she is searching correctly. What she cannot do is *arrive*, because arriving requires binding the found landmark to the running map, and that is the operation Section V says has failed. The behaviour therefore has the two properties that most distress the people who witness it: it is repetitive, because the program is centre-biased and keeps returning; and it does not stop, because the program's only stopping rule is a success signal that cannot be generated.

Grading, plainly. This is the volume's most interpretive claim and it is graded as inference. No experiment has shown that human wandering shares generative structure with insect systematic search; the transposition across phyla is enormous; and the ant's program is executed with a nervous system that has no hippocampus at all. What the comparison establishes is narrower and still worth having: that structured search following path-integration failure is a *general* solution that nervous systems implement, so that the hypothesis "the walking is a program" is a biologically ordinary hypothesis rather than a sentimental one. Section X asks whether the human data are consistent with it, and the answer is more favourable than the field's vocabulary would suggest.

X. Wandering Is Not Random — The Measurement Nobody Quotes

The word "wandering" encodes the hypothesis that the walking has no structure. That hypothesis was tested thirty-five years ago, with an unusually good method, and it failed.

The study. Forty nursing-home residents, twenty-four of them identified by nursing staff as wanderers, were observed by continuous video for thirty days, yielding more than five thousand recorded unassisted travel events, each classified by pattern (Martino-Saltzman and colleagues, 1991). This remains, by a distance, the most direct behavioural characterisation of the phenomenon.

The result. Travel was 86.8 per cent *direct* — efficient movement from an origin to a destination. Lapping, the repeated circuit of a route, accounted for 11.6 per cent. Pacing, back-and-forth movement over a short segment, accounted for 0.7 per cent. *Random travel accounted for 0.9 per cent.* And travel efficiency — the proportion of travel that was direct — was significantly related to cognitive status (Martino-Saltzman and colleagues, 1991).

What the numbers say. Three things, and each is load-bearing. First, the behaviour is overwhelmingly goal-directed; the folk category "aimless wandering" describes less than one event in a hundred. Second, the abnormality is not the presence of random travel but the *loss of directness* — cognitive decline shows up as a fall in efficiency, which is exactly what a failing position estimate produces: the trip is still aimed, it just no longer arrives. Third, the residual non-direct travel is dominated by lapping, the repeated traversal of a circuit — which is the centre-biased, returning structure that Section IX predicts and that a truly random walk does not produce.

The typology, and its later refinement. Subsequent work formalised these patterns into a wandering typology and into instruments that separate its dimensions, and reviews of the clinical literature place wandering among the most common and most consequential behavioural features of dementia, with substantial prevalence across community and institutional settings and with elopement as its feared endpoint (Cipriani and colleagues, 2014). The typology's persistence over three decades is itself evidence: a genuinely disorganised behaviour would not sort reliably into four reproducible patterns.

A caution about the inference. Directness of travel does not by itself demonstrate that the traveller has a goal in mind, and the 1991 classification was made by observers from video rather than from the traveller's report. It remains possible that direct travel reflects the constraint of corridors as much as the intent of the walker. This dissertation therefore treats the Martino-Saltzman distribution as strong evidence against the *randomness* hypothesis and as suggestive, not conclusive, evidence for the *search* hypothesis; the ledger records the distinction and

Section XVIII proposes the measurement that would settle it.

XI. Getting Lost Is Not Wandering A Distinction the Corpus Must Keep

The received account fuses two behaviours that the clinical evidence separates, and the fusion costs precision. This section separates them and then shows that one lesion can produce both.

The two behaviours. *Wandering* is repetitive locomotion within an accessible enclosure — lapping, pacing, the corridor — and it is characteristically repetitive and temporally disordered, occurring at hours when purposeful travel would not. A *missing incident* is something else. Analysis of 325 newspaper reports of persons with dementia missing in the community found that such incidents were unpredictable, non-repetitive, and — the phrase is exact and worth preserving — "temporally appropriate but spatially disordered," undertaken by multiple means of movement including car and public transport, and occurring without the discernible lapping or pacing signature of wandering. The primary antecedent was becoming lost while conducting a normal and permitted activity alone in the community (Rowe and colleagues, 2011). The authors questioned, reasonably, whether the phenomenon should be called wandering at all.

Why the distinction is not pedantic. It has a mortality attached. A review of ninety-three reported cases of persons with dementia found dead after leaving unattended found that eighty-seven per cent were discovered in natural, secluded, unpopulated areas — woods, fields, ditches, bodies of water — generally less than a mile from where they left, but often not found for a long time (Rowe and Bennett, 2003). Being found less than a mile away is the detail that should arrest the reader. This is not a person who travelled far and became lost by distance. This is a person whose positional estimate failed within a short radius of a familiar origin, and who then searched, in the wrong place, until exposure killed her. The centre-biased search of Section IX and a distance of under one mile are the same observation.

One lesion, two environments, two expressions. This dissertation proposes that the distinction between the two behaviours is environmental rather than mechanistic. Inside a bounded, familiar enclosure, an unresettable map yields a search that keeps returning to the estimated goal and is visible as repetition — wandering. Outside, on an ordinary errand in an unbounded environment, the same failure yields a single trip that does not arrive and cannot be reversed, because reversal requires a home vector that has already been corrupted — getting lost. The *temporal* appropriateness Rowe and colleagues describe is the tell: the errand was correctly scheduled and

correctly motivated, and only the spatial component failed. That is the profile of a selective navigational lesion in an otherwise intact goal-directed system.

What this costs the received account, and what it costs this one. It costs the received account the right to treat wandering and getting lost as a single phenomenon with a single management. It costs this dissertation the ease of a single behavioural target: the predictions of Section XVIII must be stated for both expressions separately, because an intervention that reduces corridor lapping is not thereby an intervention that prevents a missing incident, and the literature's own data say the two have different antecedents.

XII. The Alarm Why the Search Is Urgent

A search program explains movement. It does not explain distress, and distress is what the caregivers describe. The affective arm of wandering deserves its own mechanism, and the corpus has already characterised the structure that supplies it.

The claim to be explained. The received account posits that spatial amnesia generates intense anxiety and that the anxiety drives the search. Read literally, this makes anxiety the motor. This dissertation proposes a slightly different and more defensible relation: the search is released by the positional failure, and the noradrenergic system supplies the *gain* — the urgency, the perseveration, the resistance to redirection — rather than the initiation. A search that cannot terminate would be tolerable if it were calm. What makes it a crisis is that it is amplified.

The corpus's standing account of the noradrenergic arm. The companion volumes on the locus coeruleus establish the relevant facts and this volume imports them rather than re-deriving them: the locus coeruleus is the first nucleus in the brain to accumulate tau; its cell numbers fall while its output measures rise with severity, so the system is dysregulated rather than simply deficient; and its transmitter is simultaneously the brain's arousal, novelty and threat signal. A system whose alarm gain has risen, sitting above a navigational system that is generating a continuous unresolved error signal, is a system that will convert every uncorrected positional discrepancy into a state of urgency. The wandering that results is not calm searching; it is searching under alarm.

A pharmacological trace of the same arm. The clinical evidence that the noradrenergic system carries the affective load of these behaviours is indirect but real: prazosin, an alpha-1 adrenergic antagonist, was well tolerated and improved behavioural symptoms in patients with Alzheimer's disease with agitation and aggression in a randomised trial (Wang and colleagues,

2009). That is a demonstration about agitation, not about wandering specifically, and this dissertation does not overstate it. What it establishes is that the affective amplifier is pharmacologically reachable in a way that the navigational lesion is not — a point Section XIX takes up, because it implies that the treatable component of wandering may be the alarm rather than the search.

A structural note the corpus should record. There is an economy here of the kind the corpus has repeatedly found. The cell that governs the reset is driven by subcortical cholinergic afferents and its dysfunction in amyloid models is attributed to their loss (Leão and colleagues, 2012; Schmid and colleagues, 2016); the nucleus that supplies the alarm is the corpus's earliest tangling structure; and both are ascending neuromodulatory systems that fail early and asymmetrically. The navigational failure and the affective amplification are not two independent misfortunes that happen to coincide. They are two subcortical governors of the same cortical circuit, failing in the same window.

XIII. The Clock Why the Search Has an Hour

Wandering is not distributed uniformly across the day, and the non-uniformity is one of the most reliable observations in the whole clinical literature. An account of wandering that says nothing about when it happens is incomplete.

The observation, and its measurement. Sundowning — the late-afternoon and evening exacerbation of agitation, restlessness and wandering — was characterised against actigraphic and thermometric measures in patients with Alzheimer's disease and healthy comparison subjects. Patients showed less diurnal motor activity, a higher percentage of nocturnal activity, lower interdaily stability of motor activity, and a later activity acrophase than healthy individuals; and the severity of sundowning was associated with a later acrophase of body temperature, with lower correlation of the circadian temperature rhythm to the twenty-four-hour cycle, and with lower temperature amplitude. The data indicate that the disease disturbs circadian rhythms and that sundowning is related to a phase delay of body temperature (Volicer and colleagues, 2001).

Why a phase delay would matter to a reset. Two readings are available and this dissertation offers both, grading the second lower. The first is simply additive: a phase-delayed, low-amplitude circadian system produces an evening in which arousal is misaligned with light and with the social schedule, which raises the gain of Section XII's alarm at exactly the hour when environmental cues become least informative. The second is more specific and more speculative: the

reset requires an external reference, and the external reference degrades at dusk. The demonstration that removal of visual landmarks profoundly impairs grid periodicity and that the speed code changes in darkness (Pérez-Escobar and colleagues, 2016) is a statement about the availability of the anchor. A system that could formerly tolerate an evening because its reset was efficient becomes, once the reset is marginal, a system for which falling light is the difference between a correctable map and an uncorrectable one.

The corpus cross-link, stated and not developed. The corpus has two volumes on the chronobiological arm — one on sleep as the shared off-line restorative state across the collapse axes, one on the pineal interface and the withdrawal of melatonergic protection from the locus coeruleus. Both bear on this section and neither is re-argued here. The point this volume adds to them is narrow: if the reset is the fragile operation, then the daily hour at which the external anchor is weakest is the hour at which a marginal reset fails, and the sundowning literature's temporal signature is what that would look like from the corridor.

A caution about direction. Sundowning is associated with circadian phase delay; the association does not establish that the phase delay causes the wandering, and reverse and common-cause explanations are both available — a person who wanders at night sleeps badly, and poor sleep degrades the rhythm. The ledger grades this section as moderate for the association and as inference for the reset-specific reading.

XIV. The Cholinergic Arm One Denervation, Two Navigational Lesions

The oldest neurochemical fact in Alzheimer's disease turns out to sit on both sides of this dissertation's argument, and stating the double role resolves an apparent redundancy.

Acetylcholine drives the governor. The oriens–lacunosum-moleculare interneuron is among the hippocampal cells that receive direct subcortical cholinergic drive, and this is part of what makes it a switch rather than a fixed filter (Leão and colleagues, 2012). In an amyloid model, the mechanism identified as linking oriens–lacunosum-moleculare dysfunction to memory impairment was specifically *reduced cholinergic drive* from the septo-hippocampal pathway (Schmid and colleagues, 2016). The governor's own governor is cholinergic, and it fails.

Acetylcholine also modulates the metric. Independently, systemic blockade of muscarinic acetylcholine receptors reduces the spatial tuning of grid cells, with the effect significant

for the theta-peak-locked grid population (Newman and colleagues, 2014). The cholinergic system is therefore not only upstream of the reset; it is also a modulator of the hexagonal code itself.

The consequence for the argument. Cholinergic denervation in Alzheimer's disease is thus a *single* lesion with two distinct navigational consequences: it degrades the metric directly, and it degrades the switch that would have corrected the metric. This is not redundancy; it is the reason the navigational phenotype is so much larger than either mechanism alone would predict. It also explains, without special pleading, why cholinesterase inhibition produces a small, real, symptomatic benefit that does not alter the course: raising available acetylcholine partially supports both operations without repairing either substrate.

A reconciliation the section must state. The peptide somatostatin augments the M-current and damps repetitive firing, and muscarinic agonists antagonise the same current — the two transmitters push one conductance in opposite directions (Moore and colleagues, 1988). The companion volume used this to warn against folding the somatostatinergic deficit into the cholinergic one. That warning holds at the level of the conductance. At the level of the *circuit operation* described here, the two deficits are not opposed but concordant, because the cholinergic input's role at the oriens–lacunosum-moleculare cell is to drive the interneuron, not to act at the pyramidal conductance. One transmitter can oppose another at a channel while supporting it at a circuit. The apparent contradiction dissolves once the compartment is specified, and this dissertation states the resolution explicitly because the corpus's own earlier text invites the confusion.

XV. Measuring the Reset What Virtual Reality Detects, and What It Does Not

The diagnostic frontier the dissertation was asked to endorse is real, and it is the most successful applied consequence of the grid-cell discovery. It also has a specific and instructive limitation, which this volume regards as the more useful half of the story.

The case for the tests, stated at full strength. Standard cognitive testing leans on temporal-lobe and hippocampal function that is affected relatively late. Immersive virtual-reality path-integration testing challenges the entorhinal system directly, and the discrimination it achieves is substantially better. In forty-five patients with mild cognitive impairment and forty-one controls, patients with biomarker-positive mild cognitive impairment showed significantly larger navigation errors than biomarker-negative patients and controls; path-integration error correlated with cerebrospinal-fluid amyloid-beta and tau and with reduced entorhinal cortex volume; and the

path-integration task discriminated biomarker-positive from biomarker-negative patients with an area under the curve of 0.90, against 0.57 for the best of the standard cognitive tests (Howett and colleagues, 2019). That is not an incremental improvement. It is the difference between a test that works and a test that does not.

The signal is present decades before disease. Reduced grid-cell-like representations, altered navigational behaviour and increased hippocampal activity consistent with compensation are detectable in young adults at genetic risk (Kunz and colleagues, 2015); and path-integration deficits in risk carriers are *selective*, present when other navigational and cognitive measures are intact, and unmasked by task manipulations that remove compensatory strategies (Bierbrauer and colleagues, 2020). The lesion this dissertation describes has a behavioural read-out that precedes the clinic by an interval measured in decades.

And the deficit tracks entorhinal tau in the way the account requires. Testing 177 human volunteers on a virtual-reality task found path integration declining with age while spatial memory remained intact; and in a mouse tauopathy model, accumulation of phosphorylated tau *only in the entorhinal cortex* correlated with the navigation deficit (Koike and colleagues, 2024). The behavioural measure is not merely correlated with disease; it is correlated with pathology in the specific region the mechanism names.

The limitation, which is the section's point. All of the above establishes that virtual-reality navigation detects the *disease* early. It does not establish that it detects the *wanderer*. The same group that built much of the case tested precisely this. Virtual-reality navigation measures were related to real-world spatial disorientation in patients with Alzheimer's disease, and although allocentric wayfinding metrics significantly predicted disorientation scores within the patient group, the measures could not reliably identify which patients were at highest risk of spatial disorientation in the community (Puthusseryppady and colleagues, 2022b). Read alongside the landmark-density null of Section VIII, from the same tracking cohort, a coherent picture emerges: the tests measure the metric well and the reset not at all, and it is the reset that determines whether a person gets lost on Thursday afternoon.

What a reset-specific test would look like. The prescription follows from Section III. A path-integration test measures accumulated error at the end of an uncorrected trajectory. A reset test would measure the *correction*: it would let error accumulate, then present an unambiguous anchor — a boundary contact, a distal landmark brought into view — and measure how much of the accumulated error is discharged, and how quickly. In a healthy participant the error should collapse toward zero at the anchor. The prediction of this dissertation is that in early disease the accumulated error at the moment of anchoring will be near-normal while the *proportion discharged* will be reduced, and that the discharge fraction will discriminate real-world disorientation better than terminal path-integration error does. Section XVIII states this as a falsifiable prediction. So far as

the present search of the literature found, no existing instrument reports a discharge fraction, and the measure is available with a modest modification of tasks that already exist.

A caution on the human grid signal. Where this dissertation refers to reduced "grid-cell-like representations" in humans it means the hexadirectional functional-imaging signal (Doeller and colleagues, 2010), which is a population proxy and not a recording. Direct human grid-like recordings exist but are confined to patients with implanted electrodes (Jacobs and colleagues, 2013). No human study cited here measures grid cells in a person with Alzheimer's disease, and the volume does not imply otherwise.

XVI. The Second Correction Why a GABA-A Modulator Cannot Restore an Address

The therapeutic proposal the dissertation was asked to develop is that, because somatostatin-cell loss leaves grid cells chronically disinhibited, selective gamma-aminobutyric-acid type A receptor modulators could restore the missing inhibitory tone, re-sharpen the grid, reduce spatial anxiety and prevent wandering. The proposal is intelligent, it is the natural inference from the received chain, and this section argues that it fails at three separate points. Stating them is the volume's second deflationary contribution.

The first failure: the premise. The chain's premise is that grid cells are disinhibited by somatostatin-cell loss. Section IV disposed of this: the grid's inhibitory partner is the parvalbumin population, whose silencing degrades the hexagonal code, and somatostatin silencing does not affect grid cells at all (Miao and colleagues, 2017). A drug given to restore inhibition *to grid cells* is a drug aimed at a lesion that the earliest inhibitory casualty does not produce.

The second failure: the address. Suppose one aimed the drug at the somatostatin lesion directly, choosing a subtype whose distribution matched the affected compartment. The obvious candidate is the alpha-5 subunit, which is hippocampus-enriched and dendritically distributed and is the standard target when the intended effect is compartment-specific rather than global. The receptor literature has already localised it, and the localisation is the opposite of what the proposal requires. Alpha-5-containing receptors are preferentially targeted to the inhibitory synapses made by vasoactive-intestinal-peptide- and calretinin-positive terminals *onto the dendrites of somatostatin-expressing interneurons*, while synapses made by parvalbumin-positive inputs onto those same interneurons carry little or no alpha-5; and the receptors regulate anxiety-related behaviour through this vasoactive-intestinal-peptide-mediated phasic inhibition of interneurons (Magnin and

colleagues, 2019). The vasoactive-intestinal-peptide interneuron is the canonical *disinhibitory* cell: it inhibits other interneurons, principally somatostatin cells. A positive allosteric modulator at alpha-5 therefore amplifies the input that silences the governor. It does not restore the brake; it strengthens the hand on the brake's release.

The third failure: the direction of the field is contested, and both directions have disappointed. The pharmacology of alpha-5 in this disease has been pursued in both directions and neither has delivered a clean result. Negative modulation — the classical pro-cognitive strategy, on the reasoning that alpha-5-mediated tonic inhibition constrains hippocampal encoding — was tested with a selective negative allosteric modulator in an amyloid-precursor-protein knock-in model, and the conclusion was explicit: exposure may further compromise aberrant synapses in Alzheimer's disease, and the alpha-5 receptor is not a suitable therapeutic target for the condition (Petrache and colleagues, 2020). Positive modulation has better preclinical support — a selective alpha-5 positive allosteric modulator reverses age-related working-memory deficits and restores dendritic complexity and spine density in aged mice, with the morphological benefit persisting after treatment (Prevot and colleagues, 2021), and the same compound has been reported to reverse cognitive deficits and restore dendritic structure in both amyloid and tau transgenic models (Prevot and colleagues, 2025) — but the latter report is at published-conference-abstract level, and this dissertation grades it accordingly, as the companion volume graded the analogous somatostatin-suppression abstract.

The general principle, which is the transferable part. A positive allosteric modulator amplifies whatever gamma-aminobutyric acid is released, wherever it is released, at receptors bearing the targeted subunit. It cannot create inhibition that is not being delivered, and it cannot deliver inhibition to a compartment whose presynaptic partner is dead. When the lesion is the *loss of a cell that addressed a particular dendritic domain at a particular phase of the theta cycle*, a receptor-level drug is categorically the wrong instrument: it can change tone, and the lesion is in address and timing. This dissertation states it as a principle because it generalises beyond this receptor — the pharmacology of tone cannot supply the anatomy of address. The companion volume reached the structurally identical conclusion by a different route when it argued that restoring somatostatin *signalling* is not the same intervention as restoring somatostatin-cell *firing*, and that suppressing the cell treats the symptom of a broken governor while lowering its clearance licence further.

What is left standing, therapeutically. Two things, and they are worth naming here because they are the honest residue of the section. The first is the receptor arm the companion volume identified — somatostatin receptor subtype 1 and 4 agonism, which addresses the peptide brake and the amyloid-clearance licence and does not require the cell to be firing well (Nilsson and colleagues, 2026). The second belongs to this volume and is discussed in Section XIX: the network-hyperexcitability arm, where a drug already exists, where the trial evidence is real if partial, and where — notably for a dissertation about maps — the measured benefit included a spatial

navigation task.

XVII. The Validity Ledger

Every load-bearing joint is graded here on the corpus's three-tier convention, with the contested joints referred forward to the predictions of Section XVIII.

Established — grid cells encode a periodic spatial metric, and disrupting them impairs path integration. Directly measured: the hexagonal lattice (Hafting and colleagues, 2005); the sufficiency of recurrent inhibitory connectivity among layer II stellate cells (Couey and colleagues, 2013); and the causal demonstration that selective disruption of grid firing impairs path-integration performance without affecting other spatially selective populations (Gil and colleagues, 2018).

Established — grid representations accumulate error and are reset by the environment. Directly measured: coherent drift corrected by boundary encounters (Hardcastle and colleagues, 2015); profound impairment of grid periodicity on removal of visual landmarks, with altered speed coding in darkness (Pérez-Escobar and colleagues, 2016); experience-dependent rescaling of grids with altered geometry (Barry and colleagues, 2007). The premise on which this volume's entire argument rests is the best-established fact in it.

Established — silencing somatostatin interneurons does not degrade grid cells; silencing parvalbumin interneurons does. Directly demonstrated by pharmacogenetic silencing during recording in freely moving mice, with a clean positive-and-negative dissociation within one study, and with the anatomical basis independently measured (Miao and colleagues, 2017; Kecskés and colleagues, 2020). The volume's first correction rests on the strongest possible form of evidence for a negative claim.

Established — the somatostatin oriens–lacunosum-moleculare cell weights the entorhinal input to CA1 against the intrahippocampal input. Directly demonstrated optogenetically, with the cholinergic drive onto the interneuron demonstrated in the same work (Leão and colleagues, 2012), and with dendritic inhibition established as the dominant regulator of the CA1 pyramidal cell's input–output transformation (Lovett-Barron and colleagues, 2012).

Established — somatostatin interneurons are lost in the preclinical phase of the human disease, before excitatory neurons and before the parvalbumin class. Directly measured in the largest available multimodal human atlas, with the early-phase donors cognitively

unimpaired (Gabitto and colleagues, 2024).

Established — entorhinal layer II neurons are already profoundly depleted at the mildest clinically detectable stage. Stereological counting in human tissue: thirty-two per cent overall entorhinal loss and sixty per cent layer II loss at a Clinical Dementia Rating of 0.5, approaching ninety per cent in severe disease (Gómez-Isla and colleagues, 1996).

Established — path-integration testing detects the disease earlier and better than standard cognition. Directly measured, with biomarker anchoring and an area under the curve of 0.90 against 0.57 (Howett and colleagues, 2019); replicated in a selective deficit in risk carriers (Bierbrauer and colleagues, 2020) and related to region-specific phosphorylated tau (Koike and colleagues, 2024).

Established — wandering is overwhelmingly directed rather than random, and its abnormality is loss of efficiency. Directly measured by continuous video observation of more than five thousand travel events, with 86.8 per cent direct travel, 0.9 per cent random travel, and travel efficiency related to cognitive status (Martino-Saltzman and colleagues, 1991).

Established — missing incidents in the community are phenomenologically distinct from wandering. Directly characterised across 325 cases: unpredictable, non-repetitive, temporally appropriate but spatially disordered, arising from ordinary permitted activity, without lapping or pacing (Rowe and colleagues, 2011); with the mortality profile documented separately (Rowe and Bennett, 2003).

Established — alpha-5 GABA-A receptors are located at vasoactive-intestinal-peptide and calretinin synapses onto somatostatin interneurons, not at parvalbumin inputs to those cells. Directly measured, with the behavioural consequence tied to anxiety-related behaviour through interneuron phasic inhibition (Magnin and colleagues, 2019). This is the measured fact on which the volume's second correction rests.

Moderate — the reset is implemented by the alternation the somatostatin cell schedules. The Leão manipulation is measured; the theta-phase encoding-and-retrieval framework is a model with substantial but not decisive support (Hasselmo and colleagues, 2002); and the indirect model evidence — disrupted CA1 remapping with relatively preserved spatial responsiveness, severely lost entorhinal tuning, and impaired entorhinal-to-CA1 fast gamma coupling (Jun and colleagues, 2020) — is consistent with the account but does not manipulate the cell. The volume's signature claim is a synthesis of measured parts, and its unmeasured joint is named: no experiment shows that removing somatostatin cells abolishes boundary-driven correction of accumulated path-integration error.

Moderate — the deficit in getting lost is in the reset rather than the metric. Supported by the landmark-density null result in a real-world tracking cohort (Puthusseryppady and colleagues, 2022a), by the failure of virtual-reality navigation measures to identify who becomes disoriented in the community (Puthusseryppady and colleagues, 2022b), and by the phenomenology of the familiar-room problem. Weakened by the availability of a rival account — failure of landmark identity rather than of spatial binding — which the evidence assembled here does not exclude, and by the modest sample sizes of the tracking studies.

Moderate — cholinergic denervation degrades both the metric and the switch. Both arms are directly measured (Newman and colleagues, 2014; Schmid and colleagues, 2016; Leão and colleagues, 2012). What is not measured is their relative contribution to the human navigational phenotype, and the volume makes no quantitative claim.

Moderate — sundowning reflects circadian phase delay. The association is measured with actigraphy and thermometry (Volicer and colleagues, 2001). The direction of causation is not established, and the specific reading offered here — that failing light removes the anchor on which a marginal reset depends — is an inference layered on top of the association.

Inference — human wandering is the systematic search program. The comparative anchor is real and directly measured in its own preparation (Wehner and Srinivasan, 1981), and the human behavioural distribution is consistent with a centre-biased, non-random search (Martino-Saltzman and colleagues, 1991). But the transposition across phyla is very large; no study has analysed human wandering trajectories for the generative signature of systematic search; and directness of travel does not by itself establish an internally represented goal. This is the volume's most interpretive claim and it is graded as inference, deliberately.

Inference — the noradrenergic system supplies the urgency. The corpus's locus coeruleus volumes establish the substrate, and a randomised trial establishes that an alpha-1 antagonist improves agitation and aggression in Alzheimer's disease (Wang and colleagues, 2009). No study cited here connects noradrenergic tone to wandering specifically, and the sign of the relation in an individual patient is not established.

Inference — the somatostatin lesion's early loss removes the reserve that would have concealed the later layer II loss. This is a claim about reserve, supported by the ordering of the two clocks and by the general logic of Section III, and demonstrated by no experiment. It is the navigational analogue of the companion volume's inference that early inhibitory loss contributes causally to late excitatory loss, and it carries the same status.

Unestablished, and named as such — that any pharmacological agent can restore the reset. No compound cited in this volume is shown to restore boundary-driven error correction. The alpha-5 arm is argued here to be mis-addressed; the somatostatin receptor arm addresses the

peptide and the clearance licence rather than the timing; and the network-hyperexcitability arm acts on excitability rather than on registration. The therapeutic section is written under this constraint.

A tension the ledger records rather than resolves. Section V argues that somatostatin-cell loss disables an alternation in CA1. Section IV records that silencing somatostatin cells in the medial entorhinal cortex degrades the *aperiodic* spatial code (Miao and colleagues, 2017), while in the J20 amyloid model non-grid entorhinal spatial coding was preserved even as grid coding failed (Ying and colleagues, 2022). The companion volume recorded this anomaly and it applies here unchanged: if somatostatin cells are lost early in amyloid pathology, the aperiodic code should have suffered and did not. Acute silencing and chronic degeneration may not be equivalent; the models differ; the stages differ. None of these reconciliations is tested.

XVIII. Predictions and Falsification

Six commitments, each stated so that a definite outcome would overturn it. The first is the one that would kill the thesis.

On the reset — the load-bearing test. If the somatostatin governor schedules the correction of accumulated path-integration error, then cell-type-specific silencing or ablation of somatostatin interneurons in a navigating animal should degrade the boundary-driven *reset* of grid and place representations — measured as the residual drift remaining after a boundary contact — while leaving the moment-to-moment accumulation rate of drift, and the hexagonal periodicity itself, comparatively intact. Should somatostatin silencing leave boundary-driven error correction unimpaired, the volume's signature claim fails and the somatostatin cell has no privileged relation to the reset.

On the discharge fraction, in humans. If the reset rather than the metric determines who becomes disoriented in the world, then a virtual-reality task that permits accumulated path-integration error to be discharged at an unambiguous anchor should show, in early disease, a near-normal accumulation of error with a *reduced proportion discharged* at the anchor; and the discharge fraction should predict real-world spatial disorientation better than terminal path-integration error does. Finding the discharge fraction normal in patients whose terminal error is elevated would place the deficit in the metric after all and would remove this volume's clinical rationale.

On the search program. If wandering is systematic search rather than degraded locomotion, then trajectories recorded from wandering episodes should show the generative signature of the program — centre-biased return to a consistent estimated goal location, and loop

radius increasing over the episode — rather than the signature of a random walk or of a fixed circuit imposed by the corridor. Finding that lapping radii are constant, or that returns are to architectural features rather than to a consistent point, would reduce the program hypothesis to a metaphor and the volume would have to say so.

On the two expressions. If wandering and getting lost are one lesion in two environments, then within the same patients the severity of the reset deficit should predict both, and an environmental manipulation that supplies an unambiguous terminating anchor should reduce corridor lapping *without* reducing the rate of community missing incidents, because the missing incident occurs where no such anchor can be installed. Finding that the two behaviours dissociate on the reset measure — one predicted, the other not — would establish that they have different mechanisms and would require the account to be split.

On the pharmacology of address. If a positive allosteric modulator at alpha-5 amplifies the vasoactive-intestinal-peptide-mediated inhibition of somatostatin interneurons, then such a compound should *reduce* somatostatin-interneuron output in the hippocampus at pro-cognitive doses, measurable directly. Finding that alpha-5 positive modulation raises somatostatin-cell output, or that its pro-cognitive effect survives selective deletion of alpha-5 from somatostatin-cell dendrites, would refute this volume's address argument and would restore the receptor as a candidate for governor replacement.

On the alarm. If the noradrenergic system supplies the urgency rather than the initiation of the search, then noradrenergic antagonism should reduce the *distress and perseveration* of wandering — episode duration, resistance to redirection, associated agitation — without reducing the number of episodes initiated. Finding that noradrenergic antagonism abolishes wandering episodes altogether would make the alarm the motor rather than the amplifier and would reorder Sections IX and XII.

XIX. Corollaries for Treatment and for Care

The account yields prescriptions of two kinds, and the second kind is the one this volume regards as its practical yield, because the first is thin and honest about being thin.

What the pharmacology can currently offer, stated without inflation. Nothing restores the reset. What exists is a drug for the excitability consequence, and its evidence is real and partial. In a randomised, placebo-controlled crossover trial in Alzheimer's disease, levetiracetam did not improve the primary cognitive measure overall; but in the subgroup with epileptiform activity it improved Stroop interference naming and, notably for the present argument, significantly improved

performance on a *virtual route learning test* (Vossel and colleagues, 2021). That is a spatial-navigational benefit, in a stratified subgroup, from an anti-excitability drug — which is what one would predict if part of the navigational failure is a consequence of the network derangement that follows the inhibitory governor's loss (Palop and colleagues, 2007; Verret and colleagues, 2012). The corpus's standing caution applies: it is a subgroup result and the primary endpoint was negative. The corollary from the companion volume also stands unchanged — somatostatin receptor subtype 1 and 4 agonism remains the pharmacological arm with both loss- and gain-of-function support, and it addresses the peptide brake and the amyloid-clearance licence rather than the timing (Nilsson and colleagues, 2026).

What the pharmacology should not be asked to do. Section XVI's argument has a practical edge and it is worth putting bluntly. Sedation is not correction. A gamma-aminobutyric-acid-ergic agent given to a wandering patient will reduce the wandering by reducing the walking, and this is not the same as reducing the error that produced the search; it removes the behaviour and leaves the person inside an uncorrected map, with the fall risk that sedation carries in this population. An account that describes wandering as a search program with an unreachable success condition is an account that regards suppression of the search as the least therapeutic of the available options, and it should say so.

The environmental corollary, and its exact form. If the search terminates on binding a landmark to the running map, and if the binding operation is degraded but not absent, then the correct intervention is to make the binding as easy as the architecture can make it — not to increase the *number* of cues but to reduce the *ambiguity* of the decisive one. The wayfinding literature has developed the design principles independently of any of this mechanism, and they take exactly the form the account predicts: small, simple, legible plans; sightlines that place the destination in view from the decision point; unique and non-repeating rather than uniform corridors; and distinctive, personally meaningful cues at the goal rather than generic signage (Marquardt, 2011). The mechanistic reading this volume supplies is that a repeating corridor is pathogenic in a specific sense — it presents a system whose registration is unreliable with a scene that is genuinely ambiguous, so that even a successful binding produces the wrong answer. Repetition in the environment and repetition in the behaviour are the same failure seen from two sides.

And the caution the environmental corollary must carry. The landmark-density null result of Section VIII means this corollary cannot be stated as "more landmarks help." It did not help; environments visited by patients who became disoriented were not poorer in landmarks than those visited by patients who did not (Puthusseryppady and colleagues, 2022a). What the account predicts is narrower and testable: not density but *decisiveness* — a single unambiguous, personally salient, non-repeating anchor at the goal, visible from the point at which the search would otherwise continue. Section XVIII's fourth prediction is the trial.

The corollary for how the behaviour is described. Descriptions have consequences for care. The vocabulary of aimlessness licenses containment as a complete response; the vocabulary of search does not. The evidence in Section X is that fewer than one travel event in a hundred is random and that nearly nine in ten are direct. A person walking a corridor for the fortieth time is, on the best available measurement, executing directed travel that is failing to arrive. What that person needs is not to be stopped, and not to be sedated, but to be brought to a place her circuit can still recognise as the end of the search. The most humane implication of this dissertation is also the most mechanistically literal one: give the search something it can terminate on.

And the timing corollary, which the corpus states everywhere and must state here. The reset lesion begins in the phase in which the somatostatin cell is lost, which is the phase in which donors are cognitively unimpaired (Gabitto and colleagues, 2024). By the time wandering appears, sixty to ninety per cent of the layer II entorhinal population may also be gone (Gómez-Isla and colleagues, 1996). Every pharmacological corollary above is therefore a preclinical corollary, and every corollary that remains available at the corridor is environmental. That asymmetry is not a rhetorical flourish. It is the practical shape of the whole disease.

XX. Coda The Vector Home

There is a version of this dissertation that would have been easier to write and would have been wrong, and it is worth naming what it would have said. It would have said that the somatostatin neuron sharpens the grid; that its death blurs the hexagon; that the map dissolves; and that the person walks because the coordinate system is gone. It is a good story. It has a cell, a code, a symptom and a moral. The experiment that tests its middle says no (Miao and colleagues, 2017), and the corpus's companion volume said no before this one did.

What is left, once the easy version is discarded, is stranger and more exact. The spatial system was never accurate. It was never designed to be. It drifts, continuously, in every animal that has one, and it stays usable only because it is repaired several times a second against the world — at a wall, at a doorway, at the far corner of the room. Health is not a good map. Health is a map that can be fixed. And the fixing is a scheduled operation with an address: it happens in the thin lamina where the entorhinal cortex hands its signal to CA1, and the cell that decides when that channel is open is the cell that this disease takes first, in a phase in which the person is, by every test we have, well.

So the walking is not the map's absence. It is the map's un-repairability. A search begins because a positional estimate has expired — the same search a desert ant runs on salt pan when its

home vector runs out, loops of increasing size around the place where home should have been — and it does not stop, because the operation that would tell the system it has arrived is the operation that failed first. She sees the door. She may even recognise it. What she cannot do is fasten it to where she is. And so she walks on, and comes back, and walks on, and the loop widens, and the whole thing looks from outside like confusion and is from inside like a rescue that never completes.

The clinical record, when it has bothered to measure rather than to name, agrees. Eighty-six point eight per cent of the travel is direct; nine tenths of one per cent is random. The people who are found dead in the woods are found less than a mile from where they started, which is not the geometry of someone who wandered off but the geometry of someone who searched in the wrong place until the weather took her. Missing incidents are, in the best phrase the literature has produced, temporally appropriate and spatially disordered: the errand was right, the hour was right, the motive was right, and only the *where* had come loose.

This dissertation has declined two connections it was offered. It declines the grid, because the experiments say the hexagonal code belongs to another cell. And it declines the receptor, because the subunit the proposal would reach for sits on the governor's own dendrites, at the synapse of the cell that switches the governor off, so that the drug intended to restore the brake would press on its release. Both refusals cost the argument its neatness. Both leave it standing on measurements rather than on plausibility, which is the trade this corpus has made everywhere and should keep making.

And what remains at the end is not a molecule. It is an instruction about environments, and it is smaller and more usable than a molecule would have been. If the search cannot register its own success, then the success condition must be made unmissable: one anchor, unambiguous, unrepeated, personal, placed where the search would otherwise continue past it. Not more signs. One door that could not be any other door. That is not a treatment for Alzheimer's disease, and this volume does not pretend it is. It is the difference between a person who is walking and a person who has arrived, and for the hour that it lasts, that is the whole of it.

References

References below were retrieved and verified during preparation against the PubMed database, with the digital object identifier or PubMed identifier given as verified. Two entries are noted where they depart from that standard: reference 45 is a published conference abstract and is identified as such in the text and graded accordingly in the validity ledger, and reference

21 is a 1981 comparative-physiology paper that predates PubMed indexing of its journal and was verified against the publisher record.

1. Hafting T, Fyhn M, Molden S, Moser MB, Moser EI. Microstructure of a spatial map in the entorhinal cortex. *Nature*. 2005;436(7052):801–806. DOI: 10.1038/nature03721 (PMID: 15965463)

2. Couey JJ, Witoelar A, Zhang SJ, Zheng K, Ye J, Dunn B, Czajkowski R, Moser MB, Moser EI, Roudi Y, Witter MP. Recurrent inhibitory circuitry as a mechanism for grid formation. *Nature Neuroscience*. 2013;16(3):318–324. DOI: 10.1038/nn.3310 (PMID: 23334580)

3. Gil M, Ancau M, Schlesiger MI, Neitz A, Allen K, De Marco RJ, Monyer H. Impaired path integration in mice with disrupted grid cell firing. *Nature Neuroscience*. 2018;21(1):81–91. (PMID: 29230055)

4. Ridler T, Witton J, Phillips KG, Randall AD, Brown JT. Impaired speed encoding and grid cell periodicity in a mouse model of tauopathy. *eLife*. 2020;9:e59045. DOI: 10.7554/eLife.59045

5. Jacobs J, Weidemann CT, Miller JF, Solway A, Burke JF, Wei XX, Suthana N, Sperling MR, Sharan AD, Fried I, Kahana MJ. Direct recordings of grid-like neuronal activity in human spatial navigation. *Nature Neuroscience*. 2013;16(9):1188–1190. (PMID: 23912946)

6. Doeller CF, Barry C, Burgess N. Evidence for grid cells in a human memory network. *Nature*. 2010;463(7281):657–661. DOI: 10.1038/nature08704 (PMID: 20090680)

7. Hardcastle K, Ganguli S, Giocomo LM. Environmental boundaries as an error correction mechanism for grid cells. *Neuron*. 2015;86(3):827–839. (PMID: 25892299)

8. Pérez-Escobar JA, Kornienko O, Latuske P, Kohler L, Allen K. Visual landmarks sharpen grid cell metric and confer context specificity to neurons of the medial entorhinal cortex. *eLife*. 2016;5:e16937. DOI: 10.7554/eLife.16937 (PMID: 27449281)

9. Barry C, Hayman R, Burgess N, Jeffery KJ. Experience-dependent rescaling of entorhinal grids. *Nature Neuroscience*. 2007;10(6):682–684. (PMID: 17486102)

10. Miao C, Cao Q, Moser MB, Moser EI. Parvalbumin and somatostatin interneurons control different space-coding networks in the medial entorhinal cortex. *Cell*. 2017;171(3):507–521.e17. DOI: 10.1016/j.cell.2017.08.050

11. Kecskés M, Henn-Mike N, Agócs-Laboda Á, Szőcs S, Petykó Z, Varga C. Somatostatin expressing GABAergic interneurons in the medial entorhinal cortex preferentially inhibit layer III–V pyramidal cells. *Communications Biology*. 2020;3(1):754. DOI: 10.1038/s42003-020-01496-x

12. Leão RN, Mikulovic S, Leão KE, Munguba H, Gezelius H, Enjin A, Patra K, Eriksson A, Loew LM, Tort AB, Kullander K. OLM interneurons differentially modulate CA3 and entorhinal inputs to hippocampal CA1 neurons. *Nature Neuroscience*. 2012;15(11):1524–1530. DOI: 10.1038/nn.3235

13. Lovett-Barron M, Turi GF, Kaifosh P, Lee PH, Bolze F, Sun XH, Nicoud JF, Zemelman BV, Sternson SM, Losonczy A. Regulation of neuronal input transformations by tunable dendritic inhibition. *Nature Neuroscience*. 2012;15(3):423–430. DOI: 10.1038/nn.3024

14. Hasselmo ME, Bodelón C, Wyble BP. A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Computation*. 2002;14(4):793–817. (PMID: 11936962)

15. Nilsson P, Sörgjerd K, Kakiya N, Sasaguri H, Watamura N, Johansson L, Shimozawa M, Tsubuki S, Zhou Z, Loera-Valencia R, Takamura R, Sekiguchi M, Pegel A, Schulz S, Saito T, Iwata N, Winblad B, Saido TC.

Somatostatin receptor subtypes 1 and 4 regulate neprilysin, the major amyloid-β degrading enzyme in brain. *Journal of Alzheimer's Disease*. 2026;109(2):651–666. DOI: 10.1177/13872877251392782

16. Jun H, Bramian A, Soma S, Saito T, Saido TC, Igarashi KM. Disrupted place cell remapping and impaired grid cells in a knockin model of Alzheimer's disease. *Neuron*. 2020;107(6):1095–1112.e6. DOI: 10.1016/j.neuron.2020.06.023

17. Gabitto MI, Travaglini KJ, Rachleff VM, Kaplan ES, Long B, Ariza J, Ding Y, Mahoney JT, Dee N, Goldy J, et al. Integrated multimodal cell atlas of Alzheimer's disease. *Nature Neuroscience*. 2024;27(12):2366–2383. DOI: 10.1038/s41593-024-01774-5

18. Schmid LC, Mittag M, Poll S, Steffen J, Wagner J, Geis HR, Schwarz I, Schmidt B, Schwarz MK, Remy S, Fuhrmann M. Dysfunction of somatostatin-positive interneurons associated with memory deficits in an Alzheimer's disease model. *Neuron*. 2016;92(1):114–125. DOI: 10.1016/j.neuron.2016.08.034

19. Gómez-Isla T, Price JL, McKeel DW Jr, Morris JC, Growdon JH, Hyman BT. Profound loss of layer II entorhinal cortex neurons occurs in very mild Alzheimer's disease. *Journal of Neuroscience*. 1996;16(14):4491–4500. (PMID: 8699259)

20. Coughlan G, Laczó J, Hort J, Minihane AM, Hornberger M. Spatial navigation deficits — overlooked cognitive marker for preclinical Alzheimer disease? *Nature Reviews Neurology*. 2018;14(8):496–506. DOI: 10.1038/s41582-018-0031-x (PMID: 29980763)

21. Wehner R, Srinivasan MV. Searching behaviour of desert ants, genus *Cataglyphis* (Formicidae, Hymenoptera). *Journal of Comparative Physiology A*. 1981;142(3):315–338. DOI: 10.1007/BF00605445 *(verified against the publisher record; the journal is not indexed in PubMed for this year)*

22. Martino-Saltzman D, Blasch BB, Morris RD, McNeal LW. Travel behavior of nursing home residents perceived as wanderers and nonwanderers. *The Gerontologist*. 1991;31(5):666–672. (PMID: 1778493)

23. Cipriani G, Lucetti C, Nuti A, Danti S. Wandering and dementia. *Psychogeriatrics*. 2014;14(2):135–142. (PMID: 24661471)

24. Rowe MA, Vandeveer SS, Greenblum CA, List CN, Fernandez RM, Mixson NE, Ahn HC. Persons with dementia missing in the community: is it wandering or something unique? *BMC Geriatrics*. 2011;11:28. DOI: 10.1186/1471-2318-11-28

25. Rowe MA, Bennett V. A look at deaths occurring in persons with dementia lost in the community. *American Journal of Alzheimer's Disease and Other Dementias*. 2003;18(6):343–348. DOI: 10.1177/153331750301800612 (PMID: 14682082)

26. Puthusseryppady V, Morrissey S, Aung MH, Coughlan G, Patel M, Hornberger M. Using GPS tracking to investigate outdoor navigation patterns in patients with Alzheimer disease: cross-sectional study. *JMIR Aging*. 2022;5(2):e28222. DOI: 10.2196/28222 (PMID: 35451965)

27. Puthusseryppady V, Morrissey S, Spiers H, Patel M, Hornberger M. Predicting real world spatial disorientation in Alzheimer's disease patients using virtual reality navigation tests. *Scientific Reports*. 2022;12:13397. DOI: 10.1038/s41598-022-17634-w (PMID: 35927285)

28. Howett D, Castegnaro A, Krzywicka K, Hagman J, Marchment D, Henson R, Rio M, King JA, Burgess N, Chan D. Differentiation of mild cognitive impairment using an entorhinal cortex-based test of virtual reality navigation. *Brain*. 2019;142(6):1751–1766. DOI: 10.1093/brain/awz116 (PMID: 31121601)

29. Bierbrauer A, Kunz L, Gomes CA, Luhmann M, Deuker L, Getzmann S, Wascher E, Gajewski PD, Hengstler JG, Fernandez-Alvarez M, Atienza M, Cammisuli DM, Bonatti F, Pruneti C, Percesepe A, Bellaali Y,

Hanseeuw B, Strange BA, Cantero JL, Axmacher N. Unmasking selective path integration deficits in Alzheimer's disease risk carriers. *Science Advances*. 2020;6(35):eaba1394. DOI: 10.1126/sciadv.aba1394 (PMID: 32923622)

30. Koike R, Soeda Y, Kasai A, Fujioka Y, Ishigaki S, Yamanaka A, Takaichi Y, Chambers JK, Uchida K, Watanabe H, Takashima A. Path integration deficits are associated with phosphorylated tau accumulation in the entorhinal cortex. *Brain Communications*. 2024;6(1):fcad359. DOI: 10.1093/braincomms/fcad359 (PMID: 38347945)

31. Kunz L, Schröder TN, Lee H, Montag C, Lachmann B, Sariyska R, Reuter M, Stirnberg R, Stöcker T, Messing-Floeter PC, Fell J, Doeller CF, Axmacher N. Reduced grid-cell-like representations in adults at genetic risk for Alzheimer's disease. *Science*. 2015;350(6259):430–433. DOI: 10.1126/science.aac8128

32. Fu H, Rodriguez GA, Herman M, Emrani S, Nahmani E, Barrett G, Figueroa HY, Goldberg E, Hussaini SA, Duff KE. Tau pathology induces excitatory neuron loss, grid cell dysfunction, and spatial memory deficits reminiscent of early Alzheimer's disease. *Neuron*. 2017;93(3):533–541.e5. DOI: 10.1016/j.neuron.2016.12.023

33. Ying J, Keinath AT, Lavoie R, Vigneault E, El Mestikawy S, Brandon MP. Disruption of the grid cell network in a mouse model of early Alzheimer's disease. *Nature Communications*. 2022;13(1):886. DOI: 10.1038/s41467-022-28551-x

34. Newman EL, Climer JR, Hasselmo ME. Grid cell spatial tuning reduced following systemic muscarinic receptor blockade. *Hippocampus*. 2014;24(6):643–655. (PMID: 24493379)

35. Moore SD, Madamba SG, Joëls M, Siggins GR. Somatostatin augments the M-current in hippocampal neurons. *Science*. 1988;239(4837):278–280. DOI: 10.1126/science.2892268

36. Volicer L, Harper DG, Manning BC, Goldstein R, Satlin A. Sundowning and circadian rhythms in Alzheimer's disease. *American Journal of Psychiatry*. 2001;158(5):704–711. DOI: 10.1176/appi.ajp.158.5.704 (PMID: 11329390)

37. Wang LY, Shofer JB, Rohde K, Hart KL, Hoff DJ, McFall YH, Raskind MA, Peskind ER. Prazosin for the treatment of behavioral symptoms in patients with Alzheimer disease with agitation and aggression. *American Journal of Geriatric Psychiatry*. 2009;17(9):744–751. (PMID: 19700947)

38. Magnin E, Francavilla R, Amalyan S, Gervais E, David LS, Luo X, Topolnik L. Input-specific synaptic location and function of the α5 GABA-A receptor subunit in the mouse CA1 hippocampal neurons. *Journal of Neuroscience*. 2019;39(5):788–801. DOI: 10.1523/JNEUROSCI.0567-18.2018 (PMID: 30523065)

39. Petrache AL, Khan AA, Nicholson MW, Monaco A, Kuta-Siejkowska M, Haider S, Hilton S, Jovanovic JN, Ali AB. Selective modulation of α5 GABA-A receptors exacerbates aberrant inhibition at key hippocampal neuronal circuits in APP mouse model of Alzheimer's disease. *Frontiers in Cellular Neuroscience*. 2020;14:568194. DOI: 10.3389/fncel.2020.568194 (PMID: 33262690)

40. Prevot TD, Sumitomo A, Tomoda T, Knutson DE, Li G, Mondal P, Banasr M, Cook JM, Sibille E. Reversal of age-related neuronal atrophy by α5-GABA-A receptor positive allosteric modulation. *Cerebral Cortex*. 2021;31(2):1395–1408. DOI: 10.1093/cercor/bhaa310 (PMID: 33068001)

41. Vossel K, Ranasinghe KG, Beagle AJ, La A, Ah Pook K, Castro M, Mizuiri D, Honma SM, et al. Effect of levetiracetam on cognition in patients with Alzheimer disease with and without epileptiform activity: a randomized clinical trial. *JAMA Neurology*. 2021;78(11):1345–1354. DOI: 10.1001/jamaneurol.2021.3310 (PMID: 34570177)

42. Palop JJ, Chin J, Roberson ED, Wang J, Thwin MT, Bien-Ly N, Yoo J, Ho KO, Yu GQ, Kreitzer A, Finkbeiner S, Noebels JL, Mucke L. Aberrant excitatory neuronal activity and compensatory remodeling of inhibitory

hippocampal circuits in mouse models of Alzheimer's disease. *Neuron*. 2007;55(5):697–711. DOI: 10.1016/j.neuron.2007.07.025

43. Verret L, Mann EO, Hang GB, Barth AM, Cobos I, Ho K, Devidze N, Masliah E, Kreitzer AC, Mody I, Mucke L, Palop JJ. Inhibitory interneuron deficit links altered network activity and cognitive dysfunction in Alzheimer model. *Cell*. 2012;149(3):708–721. DOI: 10.1016/j.cell.2012.02.046

44. Marquardt G. Wayfinding for people with dementia: a review of the role of architectural design. *HERD: Health Environments Research & Design Journal*. 2011;4(2):75–90. DOI: 10.1177/193758671100400207 (PMID: 21465436)

45. Prevot TD, Bernardo AD, Wong KD, Pina-Leblanc C, Marcotte MD, Sharmin D, Mondal P, Cook JM, Sibille E. Restoring cognitive functions: unveiling the procognitive and neurotrophic impact of chronic administration of a selective α5-GABA-A receptor positive allosteric modulator in mouse models of amyloid load and tau phosphorylation. *Alzheimer's & Dementia*. 2025;20(Suppl 6):e088385. DOI: 10.1002/alz.088385 *(published conference abstract)*