THE EUROPEAN SOCIETY OF MEDICINE

Published 30 September 2026 Medical Research Archives, Volume 14 Issue 9

ORIGINAL RESEARCH

Topology is the algorithm: phrase-footprint geometry gates whether memory accumulation confers focal-lesion resilience

Fredric Narcross[1*] and Robert Marks[1]

ABSTRACT

OPEN ACCESS

The cognitive-reserve hypothesis proposes that accumulated learning protects against focal brain damage, but clinical data cannot separate the number of memories acquired from the spatial extent they occupy. We examined this question in Hexcortex, a weight-free neuromorphic architecture in which memory traces occupy identifiable locations in a hexagonal microcolumn substrate. In an angular-gyrus analog (14,160 columns), we recorded firing during a stream of two-word phrases while ablating 0–100% of the association cortical columns within a temporal memory-trace zone, using scattered and contiguous lesions (230 recordings). To separate memory count from footprint, each single-phrase substrate was retrained to approximate the other phrase's memory count while retaining its own input-zone fingerprints exactly at the byte level. This produced a balanced 2 × 2 design across eight shared lesion seeds. An initial analysis showed asymmetric drift at 75% contiguous ablation that appeared to follow memory count. The factorial analysis separated the two factors. Resilience was determined primarily by footprint rather than by the number of memories formed (7→1: $F(1,7) = 2527$, $p < 0.001$). The interaction was the largest effect at the trained transition (1→0: $F(1,7) = 944$, $p < 0.001$): additional memory improved resilience in the wide-footprint substrate ($\Delta = +0.145$, $dz = 8.9$) but not in the narrow-footprint substrate ($\Delta = -0.002$, $p = 0.19$). We describe this pattern as footprint-saturation reserve: accumulated memory improves resilience while footprint headroom remains. Footprint and phrase identity co-vary in this experiment, and the measurements concern neural patterns rather than behavioral resilience.

Keywords: cognitive reserve, neuromorphic computing, focal-lesion resilience, hippocampal–cortical memory, angular gyrus, hexagonal microcolumn architectures, lesion-symptom mapping.

AUTHOR AFFILIATIONS

[1]Department of Biotechnology Engineering, Ben-Gurion University of the Negev, Beer-Sheva, Israel.

*Corresponding author: narcross@post.bgu.ac.il

DOI

https://doi.org/10.18103/mra.2026.0590

CITATION

Narcross, F., et al., 2026. Topology is the algorithm: phrase-footprint geometry gates whether memory accumulation confers focal-lesion resilience. Medical Research Archives, [online] 14(9).

COPYRIGHT

© 2026 European Society of Medicine. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

ISSN

2375-1924

1. Introduction

Here we ask whether cognitive reserve can be examined as a substrate-level computational principle. We use Hexcortex, a weight-free neuromorphic cortical architecture in which cortical geometry selectively constrains information flow and memory traces occupy identifiable spatial locations in a hexagonal microcolumn substrate. The architecture contains a three-scale cortical hierarchy spanning phoneme, word, and phrase levels. It includes a hexagonal cone for visual pattern recognition through wavefront convergence, a dual-cone analog of the posterior superior temporal sulcus (pSTS) for polymodal stream convergence, and an angular gyrus (AG) analog for phrase-level sequential memory. Memory formation proceeds through two hippocampal–cortical projection pathways supporting categorical (Next-in-Kind, NIK) and temporal (Next-in-Time, NIT) memory, without learned synaptic weights. Computation is determined by the substrate's spatial geometry. The architecture is described in the Methods. Because phrase-level memories occupy identifiable locations, the model permits a direct test of which properties determine their survival after focal damage. The relevant clinical literature concerns cognitive reserve: individuals with similar structural brain damage can show different functional outcomes. This variation is associated with prior cognitive engagement, including education, occupational complexity, social engagement, and late-life learning (Stern, 2002, 2009, 2012; Stern et al., 2018). Reserve effects have been reported in stroke recovery, traumatic brain injury, and progressive neurodegeneration, with substantial systematic-review evidence (Cabeza et al., 2018; Valenzuela & Sachdev, 2006). Reserve is commonly distinguished from brain reserve (structural capacity) and brain maintenance (preservation of neural integrity) and refers specifically to a functional buffer between damage and deficit (Stern et al., 2018, 2020). Its substrate-level mechanism remains uncertain. Two accounts are relevant here. The accumulation account proposes that learning creates redundant memory traces that can substitute for lesioned traces. Under this account, more memories provide greater redundancy and resilience (Stern, 2009; Whalley et al., 2004). The distribution account instead relates resilience to the spatial extent of traces and their geometric relationship to lesion topography (Patterson et al., 2007; Plaut, 1995). The accounts make different predictions. Accumulation predicts an incremental benefit from each additional learned item. Distribution predicts a benefit only when learning extends or fills the spatial substrate available to the trace. Human studies have difficulty separating these variables because learning intensity usually covaries with the breadth of the representational substrate engaged (Adolphs, 2016; Bartrés-Faz & Arenaza-Urquijo, 2011). A musician with thirty years of practice has both more accumulated learning and, presumably, wider cortical engagement than one with three years. Similar covariance is expected in bilingual speakers, taxi drivers, and surgeons, where extensive learning is accompanied by broader cortical involvement (Draganski et al., 2004; Karni et al., 1998; Petersen et al., 1998). Lesion-symptom mapping can identify regions whose damage is statistically associated with deficits (Bates et al., 2003; Damasio & Damasio, 1989; Rorden & Karnath, 2004), but it cannot readily determine whether memory count or representational breadth before injury provides the buffer. Computational substrates offer a way to manipulate count and spatial spread independently while keeping memory traces in identifiable locations (Adolphs, 2016; Plaut, 1995).

Computational connectionist neuropsychology has long used this approach. Plaut showed that random lesions of distributed representations could reproduce double dissociations previously interpreted as evidence for modular cognitive architectures (Plaut, 1995). Later work emphasized that the spatial distribution of representations, rather than the number of memories alone, predicts which lesions permit functional recovery (McClelland et al., 1995; Patterson et al., 2007). Hexcortex extends this approach because phrase-transition memories in the AG analog are stored in identifiable association cortical columns. The input-fingerprint

footprint of each phrase—the set of columns in the input-axis zone (IAZ) activated by the cascade from the phrase's pSTS analog output—is discrete, measurable, and experimentally accessible.

The study had two phases. Phase 1 established a baseline dose–response by recording continuous input to the angular gyrus substrate (Seghier, 2013) (Fig. 1) while ablating 0–100% of the association cortical columns in its memory-trace zone (the NIT zone, NTZ). Lesions were either random (scattered) or contiguous (focal). At 75% contiguous lesion, substrates with more accumulated memories showed less drift from the canonical (unlesioned) firing pattern. This result initially appeared consistent with the accumulation account. However, substrates that differed in memory count had also been trained on different phrases and therefore differed in the spatial footprint occupied by those phrases. Phase 2 used matched-count derivative substrates, each trained on one phrase to a memory count approximating that of the original substrate trained on the other phrase. The resulting 2 × 2 design varied footprint and memory count factorially. The phase-1 asymmetry was associated with the spatial footprint of the trained phrase's input fingerprint, which determines the geometric distribution of memory traces in the NTZ, rather than with memory count alone. Additional training improved resilience in the substrate whose footprint allowed further growth, but not in the substrate whose footprint was already dense relative to the available area. This dependence appeared as a footprint × count interaction, which we call footprint-saturation reserve. Two constraints apply to the interpretation. First, the vocabulary contains only two phrases, so footprint width and phrase identity are fully confounded; footprint is therefore an interpretation supported by converging evidence rather than an independently manipulated variable. Second, Hexcortex is a model. Its correspondences with clinical cognitive reserve (Cabeza et al., 2018; Stern, 2009, 2012; Stern et al., 2018, 2020) are predictions of the architecture, not demonstrated properties of biological cortex.

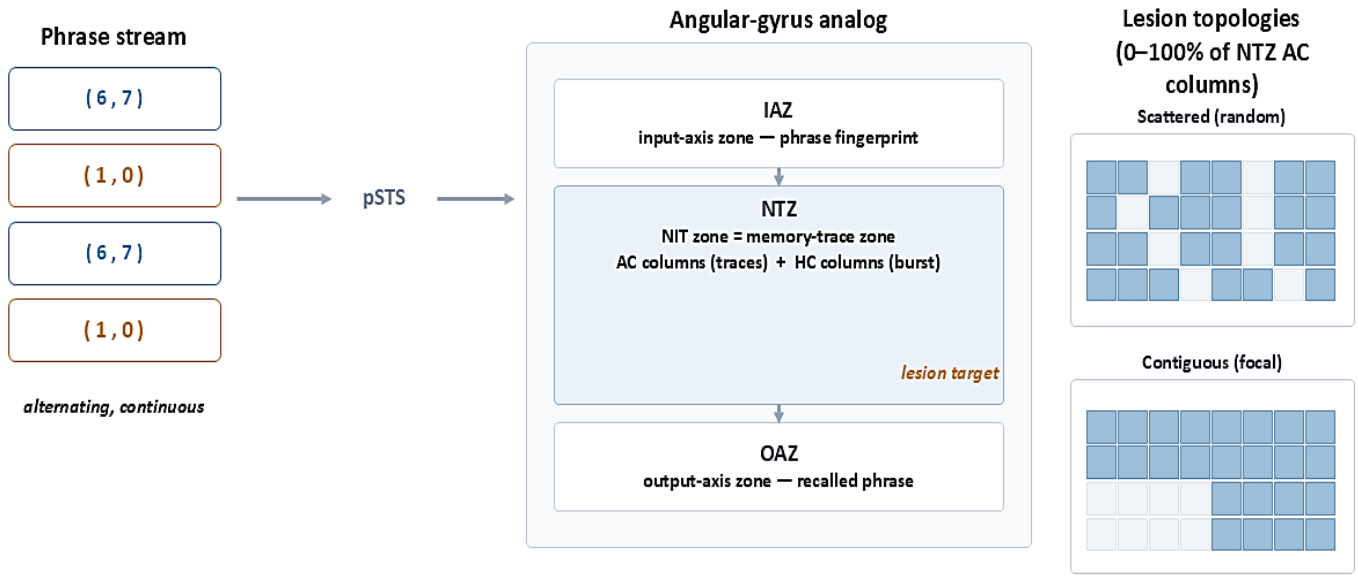

Fig. 1 | Experimental design: focal lesions in the angular-gyrus analog. The angular-gyrus analog receives a continuous stream of alternating two-word phrases ((6,7) and (1,0)) via the posterior superior temporal sulcus (pSTS). Phrase fingerprints enter the input-axis zone (IAZ); next-in-time (NIT) memories are stored on association cortical (AC) columns in the NIT memory-trace zone (NTZ) and consolidated by hippocampal cortical column (HC) bursts; the recalled phrase is read out at the output-axis zone (OAZ). Phase 1 ablates 0–100% of the NTZ AC columns under two lesion topologies: scattered (randomly distributed) and contiguous (a spatially connected block).

2. Methods

HEXCORTEX ARCHITECTURE

Hexcortex is a weight-free neuromorphic cortical architecture (Davies et al., 2018; Roy et al., 2019) in which information processing and memory formation are carried by the spatial geometry of a hexagonal microcolumn substrate rather than by learned synaptic weights. The architecture comprises a three-scale cortical hierarchy (Felleman & Van Essen, 1991): a source hexagonal cone that implements visual pattern recognition through wavefront convergence; a dual-cone analog of the posterior superior temporal sulcus that implements polymodal stream convergence; and an angular gyrus analog that implements phrase-level sequential memory. Within each grid, microcolumns are arranged on a contiguous hexagonal lattice with each column connected to its six immediate neighbors; activation cascades propagate through the lattice along trajectories determined entirely by the substrate geometry and the structural connectivity of the participating columns. Cascade activity is driven by Hz-derived oscillatory cycles (gamma 40 Hz at the source grid, theta 6 Hz at pSTS, delta 1.5 Hz at AG). It is gated by hippocampal–cortical interactions that support two distinct memory-formation pathways: NIK, supporting categorical memory, and NIT, supporting temporal memory.

The substrate of interest in the present work is the angular-gyrus (AG) analog, a hexagonal microcolumn grid of 14,160 columns arranged in a 120 × 118 anterior–posterior × dorsal–ventral mesh (Markram et al., 2015; Mountcastle, 1997). Columns are typed as either association cortical columns (AC; 10,883 columns, 76.9%) or hippocampal cortical columns (HC; 3,277 columns, 23.1%). Fascicular input and output to and from the AC columns is mediated by association fascicular columns (AFC), with fascicular efferents originating in Layer II and fascicular afferents terminating in Layer III. HC columns are distributed across three dorsal–ventral tiers (dorsal: 1,600; middle: 1,053; ventral: 624) at densities inversely proportional to their dorsoventral position, modeling the gradient of hippocampal–cortical input observed in the primate angular gyrus (Seghier, 2013). The substrate is functionally partitioned into three zones along the anterior–posterior axis (Fig. 1): an input-axis zone (IAZ, columns 0–19, 2,360 columns), a NIT zone (NTZ, columns 20–99, 9,440 columns), and an output-axis zone (OAZ, columns 100–119, 2,360 columns).

NTZ is the substrate for NIT memory formation: pairs of activated AC columns that co-fire under HC burst become bound through a hippocampal-mediated potentiation rule (Buzsáki, 2002; Eichenbaum, 2000), after which subsequent activation of one member of the pair can recall the other. NIT memories therefore occupy identifiable spatial locations within the NTZ, and the spatial extent and distribution of memories formed by training on a particular input stream can be observed directly. This property, that memory traces are stored in spatially addressable locations rather than distributed across learned weights, is what makes the architecture suitable for the substrate-level lesion-resilience analyses reported below.

PHRASE VOCABULARY AND IAZ FINGERPRINTS

The training vocabulary consisted of two phrases composed of four digits: phrase (6,7) and phrase (1,0). Visual fingerprints for each digit were derived from the UCI Optical Recognition of Handwritten Digits dataset (Alpaydin & Kaynak, 1998) (8×8 grayscale, accessed via scikit-learn's load_digits function) using a deterministic top-k selection of strong-gradient pixels. Auditory fingerprints were derived from the AudioMNIST dataset (Becker et al., 2024) (speaker 01, 5 recordings per digit, digits 0, 1, 6, 7, at 48 kHz) using mel-spectrogram analysis with peak-frame top-k selection per theta cycle (Buzsáki, 2002; Buzsáki & Wang, 2012; Colgin, 2013; Lisman & Jensen, 2013). Vocabulary fingerprints drive cascade activity through the pSTS substrate (a separate

dual-cone hexagonal grid of 5,884 microcolumns spanning 79 AP × 120 DV, implementing audiovisual integration), whose multimodal meeting zone (MMZ) output is mapped to the angular-gyrus IAZ via a surgical projection protocol producing per-digit IAZ activation patterns. Four substrate conditions were used: canonical (both phrases), only67 (phrase (6,7) only), only10 (phrase (1,0) only), and hclesion50 (both phrases, encoded under 50% random HC ablation). For the substrates used in this study, the IAZ fingerprint sizes were: for canonical, d6 = 320, d7 = 320, d1 = 440, d0 = 400 cells; for hclesion50, d6 = 320, d7 = 320, d1 = 460, d0 = 460 cells; for only67, d6 = 300, d7 = 320, d1 = 440, d0 = 440 cells; for only10, d6 = 300, d7 = 320, d1 = 460, d0 = 420 cells. Variability across substrate conditions reflects stochastic sample selection in the AudioMNIST audio extraction during cache construction. Within a substrate condition, fingerprints are deterministic once cached.

AG NIT TRAINING

NIT training operates on phrase pairs. For a phrase (n, n+1), the AG is presented with digit-n's IAZ fingerprint as input for 60 simulation ticks (with NIT memory formation frozen), followed by a word-transition operation that swaps the input to digit-(n+1)'s fingerprint, followed by a delta window during which NIT memory formation is enabled. Co-presence pairs identified in this delta window are bound as new NIT memories. After each phrase presentation, a 20-tick refractory phase resets the cascade state with frozen memory. For canonical substrates, both phrases are trained in an interleaved manner over 10 cycles. For phrase-isolated substrates (only67, only10), only the named phrase is trained over the same number of cycles. For HC-lesion substrates (hclesion50), both phrases are trained, but with a random subset of HC columns silenced during training, modeling anterograde-amnesia-analog damage during encoding (Frankland & Bontempi, 2005; Scoville & Milner, 1957; Squire, 1992).

MATCHED-COUNT DERIVATIVE SUBSTRATES

To construct a matched-count derivative substrate (e.g., only67_matched at target memory count 934), we loaded the IAZ fingerprints from the source cache (only67 in this example, preserving byte-identity of d6/d7/d1/d0 patterns), reset the AG NIT memory state, and ran the cycle-by-cycle training loop for the target phrase only, snapshotting the AG memory state after each cycle. After the loop completed, the snapshot whose memory count was closest to the target was restored and saved as a new derivative cache, with metadata recording its lineage from the source cache and the target count. Because (a) the source-cache fingerprints are inherited byte-identically, (b) the AG engine is fully deterministic given fixed inputs and starting state, and (c) the AG reset operation produces a known clean initial state, the rebuild protocol guarantees that the matched substrate differs from a hypothetical 'from-scratch' rebuild only in lineage metadata; the actual memory state and fingerprint contents are reproducible. We verified this by rebuilding only67 with target = 532 (the original count); the rebuild produced an AG memory count of exactly 532 at cycle 10. For the experiment reported here, the two matched substrates achieved: only67_matched, target 934, achieved 926 at cycle 17 (gap = −8, 0.9% short); only10_matched, target 532, achieved 556 at cycle 6 (gap = +24, 4.5% over). The discrete granularity of cycle counts limits the match precision: the (1,0) phrase grows at ~108 memories per cycle, so cycle 5 produces 448 memories and cycle 6 produces 556; no integer-cycle landing produces 532 exactly. We accept this granularity rather than attempting fractional-cycle training, which would introduce protocol complications absent in the original substrate construction.

LESION PROTOCOLS

NTZ_AC lesion masks are constructed by enumerating all 6,163 columns in the NTZ association cortical pool, applying a permutation seeded by a fixed integer (--lesion-seed), and silencing the first round (fraction ×

6163) columns under that permutation. For random lesion topology, the permutation is uniform over the full NTZ_AC pool; silenced columns are scattered throughout the substrate. For contiguous lesion topology, the ordering is structured to produce focal silencing: a seed-selected column of the NTZ_AC pool is taken as the center, the remaining pool columns are ranked by Euclidean distance from it in anterior–posterior by dorsal–ventral coordinates, and the nearest columns are silenced until the target count is reached. The result is a compact, spatially concentrated lesion of the requested size. Critically, lesion permutations for a given seed are nested across fractions: the 75% mask is a strict superset of the 50% mask, which is a strict superset of the 25% mask. This allows 'more damage' to be characterized as a clean superset relation across fractions rather than independent random samples, mirroring the progression observed in clinical degeneration studies (Bates et al., 2003; Rorden & Karnath, 2004; Stam, 2014).

CONTINUOUS-INPUT RECORDING

For each substrate × lesion mode × seed combination, the AG was loaded from its cache, the lesion mask applied, and the substrate driven by a phrase-alternating word stream: digits [6, 7, 1, 0, 6, 7, ...] with 120 ticks per word, for 60 total words (7,200 ticks ≈ 40 s simulated time per recording). The stream contains four transition types: trained transitions for the substrate's phrase (e.g., 6→7 for only67); cross-phrase transitions (7→1 and 0→6) that are not directly trained in any substrate; and skipped transitions (e.g., 1→0 for only67) that are within the canonical training set but were not trained in the phrase-isolated substrate. During recording, the per-tick NTZ AC firing pattern is logged with a downsampling factor of 2 (every second tick), producing 3,600 sample frames per recording. Per-transition mean firing patterns are computed by averaging frames within each word's transition window. Recording metadata records the lesion mask, seed, fraction, substrate cache, and ground-truth transition labels for downstream analysis.

DEFICIT-FINGERPRINT COSINE ANALYSIS

For each non-canonical recording, the canonical recording with matching lesion mode, fraction, and seed is identified (apples-to-apples masking; both recordings use the same column-silencing mask). Per-transition, mean NTZ AC firing patterns are computed for both the target recording and the canonical match, and the cosine similarity between them is computed as the inner product of the two vectors divided by the product of their Euclidean norms (Diedrichsen & Kriegeskorte, 2017; Kriegeskorte et al., 2008). Per-transition cosines are aggregated across the 8 seeds per cell to produce mean ± standard deviation summaries. The reserve experiment 2 × 2 reads four cell values per transition; the count axis is characterized as the within-phrase mean absolute difference, and the phrase axis as the within-count mean absolute difference.

STATISTICAL ANALYSIS

Each transition in the matched-count factorial was analyzed by two-way repeated-measures analysis of variance with footprint (phrase family) and memory count as within-subject factors and lesion seed as the subject term (n = 8 seeds, df_error = 7). The repeated-measures form is the correct one for this design because the same eight lesion masks were applied to every substrate, so each seed contributes a complete 2 × 2 and the between-substrate comparisons are paired rather than independent. Effect sizes are reported as partial eta squared, computed as F/(F + df_error). Planned simple effects, the count contrast within each footprint family and the footprint contrast at each count level, were tested with paired t-tests on the same seed and are reported with 95% confidence intervals on the mean difference and Cohen's dz. No correction for multiple comparisons across the four transitions was applied; the two transitions on which the argument rests are significant at magnitudes that survive any conventional correction. Eight seeds is a small sample by behavioral

standards, but the engine is deterministic and the design fully paired, so the seed term captures variation in lesion placement rather than measurement noise; the contrasts on which the argument rests carry paired effect sizes of dz = 6.0 to 8.9, well outside the range in which sample size limits inference. Analyses were performed in Python (SciPy, statsmodels).

3. Results

DOSE-RESPONSE ACROSS SUBSTRATES AND LESION TOPOLOGIES

We measured continuous-input AG recordings under four substrate conditions (Table 1, Fig. 2): canonical, trained on both digit phrases (1,358 NIT memories after 10 cycles); only67, trained on the (6,7) phrase only (532 memories after 10 cycles); only10, trained on the (1,0) phrase only (934 memories after 10 cycles); and hclesion50, trained on both phrases under 50% random hippocampal cortical column (HC) ablation during encoding (278 memories), modeling anterograde-amnesia-analog damage during encoding (Scoville & Milner, 1957; Squire, 1992). For each substrate, we recorded the AG response to a 60-word phrase-alternating stream under lesions of 0%, 25%, 50%, 75%, and 100% of the NTZ association cortical column (NTZ_AC) population, applied in both random (scattered) and contiguous (focal) topologies, across 8 random seeds per non-zero fraction (n = 230 recordings total in phase 1).

Table 1 | Substrate conditions used in phase 1

Substrate	Trained phrases	Training cycles	Encoding lesion	NIT memory count
canonical	(6,7) and (1,0)	10	none	1,358
only67	(6,7)	10	none	532
only10	(1,0)	10	none	934
hclesion50	(6,7) and (1,0)	10	50% random HC ablation	278

The deficit-fingerprint metric, per-transition cosine similarity (Diedrichsen & Kriegeskorte, 2017; Kriegeskorte et al., 2008) of the lesioned substrate's NTZ AC firing pattern to the canonical substrate's firing pattern under a matched lesion mask, produced expected dose–response curves under random lesion topology: similarity decreased monotonically from above 0.90 at 0% to undefined at 100% ablation (firing collapses to zero), with comparable trajectories across substrate conditions. Because random ablation removes columns from spatially unbiased locations, it degrades the trace substrate evenly rather than concentrating damage in any one region; each substrate therefore loses a comparable share of the columns supporting a given trace, and the per-fraction noise structure was similar across substrates.

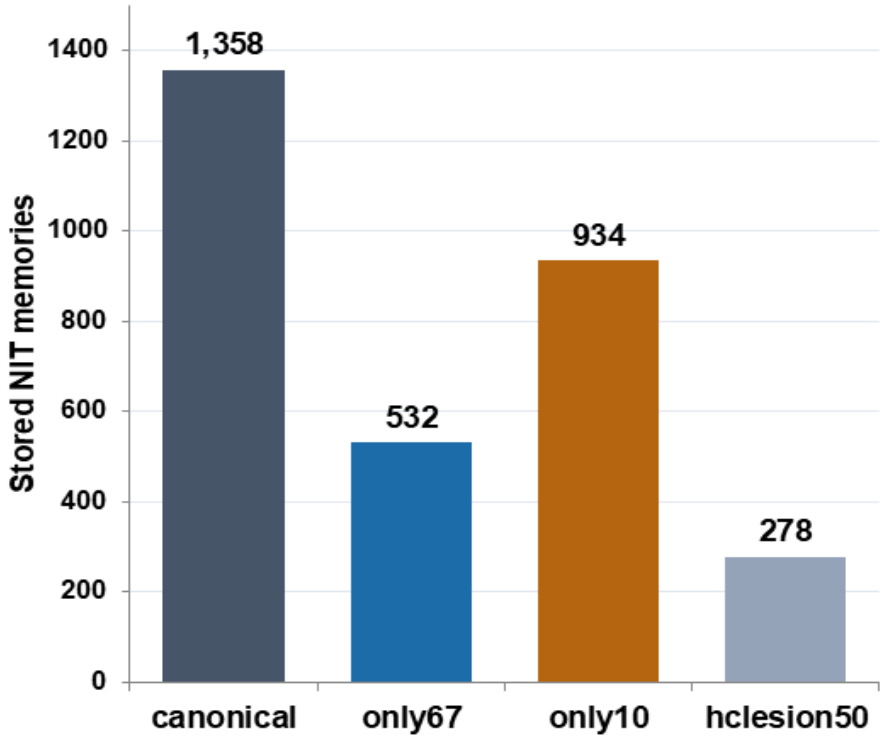

Fig. 2 | Baseline substrate conditions and stored-memory counts. The four phase-1 substrates: canonical encodes both digit phrases; only67 and only10 each encode a single phrase; hclesion50 encodes both phrases under 50% random hippocampal cortical column (HC) ablation during encoding. Bars show the resulting stored NIT memory counts for each substrate (1,358, 532, 934, and 278, respectively, after 10 training cycles).

Contiguous lesion topology produced a markedly different pattern. While trained transitions in the (6,7)-isolated substrate (only67) maintained similarity above 0.81 even at 75% contiguous lesion, the corresponding cross-phrase transitions (which had no trained NIT memory) showed substantial drift; for example, the 1→0 similarity in only67 dropped from 0.45 at 50% lesion to 0.27 at 75%. The asymmetry between same-substrate trained and skipped transitions is mechanically expected: skipped transitions rely entirely on cascade-only activity through unmemorized regions, which is more vulnerable to focal damage than trace-memorized activity.

At 75% contiguous lesion, a further between-substrate asymmetry appeared. The substrate with more memories (only10, 934) drifted less from the canonical pattern than the substrate with fewer memories (only67, 532). Mean recording file size showed the same ordering: 7.9 ± 0.4 MB for only67, 9.6 ± 0.5 MB for only10, and 8.4 ± 0.6 MB for hclesion50 (278 memories). We report file size only as a descriptive feature of the recording pipeline because it depends on compression and is not a calibrated measure of firing. Together with the cosine results, the phase-1 pattern appeared to support a memory-count dose–response in which additional NIT memories provided additional spatial redundancy (Stern, 2009; Whalley et al., 2004). That interpretation could not be tested in phase 1 because memory count was confounded with phrase identity and, consequently, with input-fingerprint footprint. Phase 1, therefore, could not distinguish the effects of count, footprint, or their combination.

MATCHED-COUNT SUBSTRATE CONSTRUCTION

To separate these variables, we built two derivative substrates (Fig. 3): only67_matched, trained on the (6,7) phrase to 926 NIT memories (17 cycles) while inheriting only67's IAZ fingerprints byte-identically; and only10_matched, trained on the (1,0) phrase to 556 NIT memories (6 cycles) while inheriting only10's IAZ fingerprints byte-identically. Construction began by loading the source cache's phrase-specific IAZ fingerprints, resetting the NIT memory state, and repeating the cycle-by-cycle training loop for the selected phrase until the closest available memory count was reached. The fingerprints were loaded from disk rather than regenerated. Within each phrase family, the matched substrate therefore differs from its source in NIT memory count and number of training cycles, while the other measured inputs are held fixed.

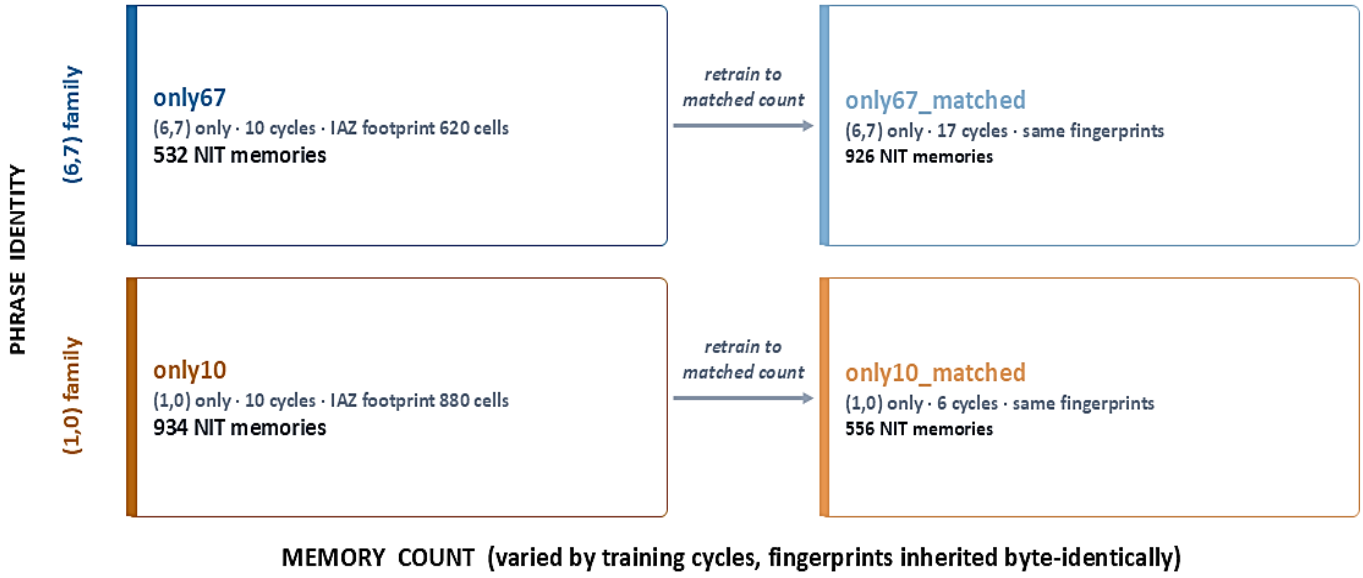

Fig. 3 | Matched-count substrate construction. (a) Per-cycle memory-formation trajectory: AG NIT memory count versus training cycle for (6,7)-only and (1,0)-only training, each using its own source-cache IAZ fingerprints; both grow approximately linearly, with target counts and best-match cycles marked. (b) Construction of the matched-count derivatives: each single-phrase substrate is retrained to a memory count characteristic of the other phrase while inheriting its own IAZ fingerprints byte-identically, factorially crossing phrase identity with memory count.

The training trajectories also differed by phrase (Fig. 3). The (6,7) phrase produced approximately 56 new NIT memories per training cycle across 22 cycles and did not reach growth saturation over the range examined. The (1,0) phrase produced approximately 108 new memories per cycle, nearly twice the rate. This difference follows the IAZ fingerprint sizes: the (6,7) fingerprints span 620 NTZ cells (d6: 300; d7: 320), whereas the (1,0) fingerprints span 880 cells (d1: 460; d0: 420). Larger fingerprints generate wider IAZ→NTZ cascades, increasing the number of co-presence events per HC burst (Buzsáki, 2002; Lisman & Jensen, 2013) and therefore the number of NIT memories formed per cycle.

FOOTPRINT GEOMETRY GATES THE RESILIENCE BENEFIT OF ADDED MEMORY

We applied the focal-lesion protocol (75% contiguous, seeds 0–7) to the two derivative substrates and computed the deficit-fingerprint cosine-to-canonical for each transition in the recorded stream (Table 2, Fig. 4). Each transition was analyzed by repeated-measures ANOVA with footprint (phrase family) and memory count as within-subject factors (Table 3). The footprint factor was dominant for the transitions measured with sufficient precision, but the interaction between footprint and count was the main factorial result. For the transition

most heavily trained in each family (6→7 for the only67 family and 1→0 for the only10 family), the (6,7)-trained substrates scored 0.816 ± 0.152 (only67, 532 memories) and 0.803 ± 0.141 (only67_matched, 926 memories) at 6→7. The (1,0)-trained substrates scored 0.702 ± 0.086 (only10_matched, 556) and 0.702 ± 0.087 (only10, 934). Although the ordering was as expected, no factor reached significance at this transition (footprint $F(1,7) = 2.50$, $p = 0.158$; count $F(1,7) = 1.60$, $p = 0.246$; interaction $F(1,7) = 1.75$, $p = 0.228$). The paired contrasts were also weak (footprint at low count $\Delta = 0.115$, 95% CI −0.054 to 0.283, $p = 0.152$; at high count $\Delta = 0.101$, 95% CI −0.054 to 0.255, $p = 0.166$). Seed-to-seed variance at 6→7 was large and shared across substrates. So, the main interpretation relies on the two transitions measured more precisely.

Table 2 | Cosine similarity to canonical at 75% contiguous lesion (n = 8 seeds, mean ± SD). Each transition is shown across the 2 × 2 of phrase identity and memory count.

Transition	only67 (532)	only67_matched (926)	only10_matched (556)	only10 (934)
6→7 (trained / skipped)	0.816 ± 0.152	0.803 ± 0.141	0.702 ± 0.086	0.702 ± 0.087
7→1 (cross-phrase)	0.705 ± 0.020	0.702 ± 0.021	0.860 ± 0.015	0.929 ± 0.013
1→0 (skipped / trained)	0.273 ± 0.130	0.271 ± 0.128	0.768 ± 0.013	0.913 ± 0.009
0→6 (cross-phrase)	0.579 ± 0.282	0.571 ± 0.274	0.603 ± 0.050	0.594 ± 0.057

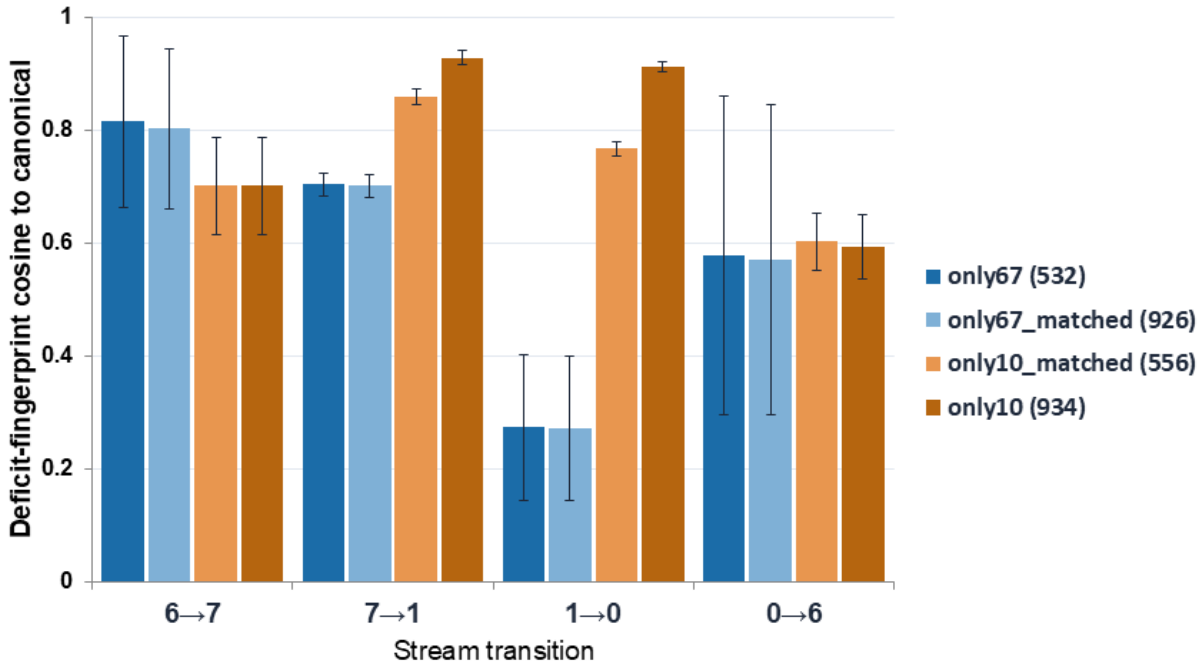

Fig. 4 | Lesioned firing pattern tracks footprint family rather than memory count. Deficit-fingerprint cosine similarity to the canonical (unlesioned) firing pattern under 75% contiguous lesion, by stream transition, for the four matched-count substrates (Table 2; mean ± s.d. over 8 seeds). Within a footprint family, the low-count and matched-count bars are near-identical, whereas across families they diverge; the divergence is resolved with high precision at 7→1 and 1→0, is large but not statistically reliable at 6→7, and is unresolvable at 0→6, where seed variance is greatest.

At the 1→0 transition, the difference between phrase families was larger and more precisely measured. The (6,7) family remained near 0.27 at both counts (0.273 ± 0.130 and 0.271 ± 0.128), whereas the (1,0) family scored 0.768 ± 0.013 at 556 memories and 0.913 ± 0.009 at 934. The footprint main effect was $F(1,7) = 167.1$, $p < 0.001$, $\eta p^2 = 0.96$. The paired footprint contrast was $\Delta = -0.495$ (95% CI −0.593 to −0.396) at low count and $\Delta = -0.642$ (95% CI −0.752 to −0.533) at high count. The cross-phrase 7→1 transition showed the same ordering with the highest measurement precision of the four transitions (footprint $F(1,7) = 2526.9$, $p < 0.001$, $\eta p^2 = 1.00$). At 0→6, no factor approached significance (all $p > 0.11$), and per-cell standard deviations reached 0.28. We therefore treat this transition as uninformative rather than as evidence for a null effect. For the adequately resolved transitions, footprint accounted for nearly all between-substrate variance. The phase-1 asymmetry initially associated with memory count was thus explained mainly by the spatial footprint of the trained representation. This is consistent with predictions from connectionist neuropsychology (Patterson et al., 2007; Plaut, 1995), but here the effect is isolated in a substrate with explicit memory-trace localization.

Table 3 | Two-way repeated-measures ANOVA of cosine similarity to canonical at 75% contiguous lesion, by transition. Factors are footprint (phrase family) and memory count, with lesion seed as the subject term; F ratios are on (1, 7) degrees of freedom, n = 8 seeds. ηp^2 is partial eta squared. The 6→7 and 0→6 transitions are measured with insufficient precision to resolve any effect.

Transition	Footprint main effect	Memory-count main effect	Footprint × count interaction
6→7	$F = 2.50$, $p = 0.158$, $\eta p^2 = 0.26$	$F = 1.60$, $p = 0.246$, $\eta p^2 = 0.19$	$F = 1.75$, $p = 0.228$, $\eta p^2 = 0.20$
7→1	$F = 2526.9$, $p < 0.001$, $\eta p^2 = 1.00$	$F = 205.1$, $p < 0.001$, $\eta p^2 = 0.97$	$F = 136.5$, $p < 0.001$, $\eta p^2 = 0.95$
1→0	$F = 167.1$, $p < 0.001$, $\eta p^2 = 0.96$	$F = 407.5$, $p < 0.001$, $\eta p^2 = 0.98$	$F = 944.1$, $p < 0.001$, $\eta p^2 = 0.99$
0→6	$F = 0.04$, $p = 0.841$, $\eta p^2 = 0.01$	$F = 3.25$, $p = 0.115$, $\eta p^2 = 0.32$	$F = 0.01$, $p = 0.922$, $\eta p^2 = 0.00$

MEMORY-COUNT EFFECTS DEPEND ON FOOTPRINT HEADROOM

The footprint effect coexisted with a count effect, but the count effect depended on phrase footprint. The footprint × count interaction was large at both well-resolved transitions (1→0: $F(1,7) = 944.1$, $p < 0.001$, $\eta p^2 = 0.99$; 7→1: $F(1,7) = 136.5$, $p < 0.001$, $\eta p^2 = 0.95$), with the benefit of additional memories confined to the wide-footprint family. At 1→0, the interaction exceeded both main effects; at the cross-phrase 7→1 transition, the footprint main effect was larger. This difference is consistent with the functions engaged by the two transitions. Stored memories contribute directly at 1→0, whereas the 7→1 pattern is produced by cascade activity associated with footprint geometry. At 1→0, additional training increased similarity in the (1,0) family from 0.768 ± 0.013 to 0.913 ± 0.009 ($\Delta = +0.145$, 95% CI +0.131 to +0.159, $t(7) = 25.0$, $p < 0.001$, $dz = 8.9$). The same manipulation had no measurable effect in the (6,7) family ($\Delta = -0.002$, 95% CI −0.006 to +0.002,

p = 0.193). The 7→1 transition showed a similar pattern (wide family Δ = +0.069, 95% CI +0.059 to +0.079, t(7) = 16.9, p < 0.001, dz = 6.0; narrow family Δ = −0.004, p = 0.341). Memory count therefore contributed to resilience only in the family with remaining footprint capacity.

The interaction can be interpreted in terms of fingerprint density (Fig. 5). The (6,7) phrase has a 620-cell IAZ fingerprint footprint within which its NIT memories form. At 532 memories, the substrate contains 0.86 memories per fingerprint cell; at 926 memories, the density rises to 1.49. The results indicate that this footprint was functionally saturated by 10 training cycles, because further cycles added memories without a measurable increase in resilience. The (1,0) phrase has an 880-cell footprint. Its density is 0.63 memories per cell at 556 memories and 1.06 at 934 memories. At 556 memories, the (1,0) substrate remains below the density already reached by the (6,7) substrate after 10 cycles, and additional training continues to improve resilience. By 934 memories, its density is close to that of the (6,7) substrate at 532 memories and may be approaching saturation. The interaction distinguishes this account from a simple claim that memory count has no effect. It also resembles the capacity limit of Hopfield-style associative memory networks, where retrieval saturates when the ratio of stored patterns to substrate size crosses a critical threshold (Hopfield, 1982). In the present model, the corresponding quantity is fingerprint density rather than the pattern-to-capacity ratio. Because only two footprint widths were tested, the saturation threshold is inferred rather than mapped directly.

We refer to this pattern as *footprint-saturation reserve*. Additional training contributes to resilience according to memory density relative to the fingerprint footprint, rather than absolute memory count alone. Within-family count effects appear below saturation density and are absent above it.

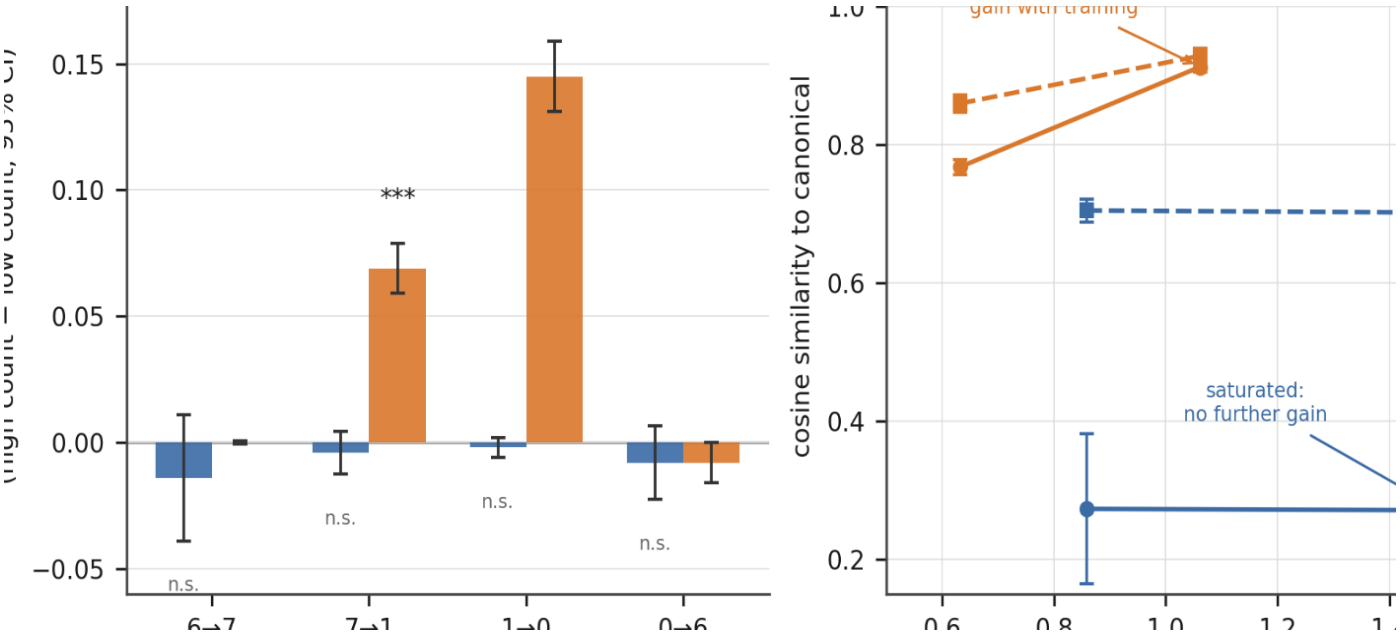

Fig. 5 | Footprint-saturation interpretation. Count effect (memory count varied, footprint fixed) shown separately within each footprint family, from the matched-count factorial (Table 3). Additional training raises similarity in the wide-footprint (1,0) family (1→0: Δ = +0.145, p < 0.001) but not in the narrow-footprint (6,7) family (Δ = −0.002, p = 0.976), the interaction that defines the account: resilience is gated by trace density relative to the phrase's IAZ fingerprint footprint ((6,7): 620 NTZ cells; (1,0): 880 NTZ cells), so accumulated memory contributes only below saturation.

CROSS-PHRASE HALO ASYMMETRY

The cross-phrase transitions (7→1 and 0→6) were not directly trained in any substrate and might therefore be expected to behave similarly across substrate conditions. Instead, cosine similarity to the canonical pattern was consistently higher for the (1,0)-trained substrates than for the (6,7)-trained substrates. At 7→1, the values were 0.860 ± 0.015 (only10_matched), 0.929 ± 0.013 (only10), 0.705 ± 0.020 (only67), and 0.702 ± 0.021 (only67_matched). This difference is consistent with the input-fingerprint footprints. The wider (1,0) fingerprints generate wider NTZ cascade activity even when the transition does not directly engage a trained NIT memory. The resulting halo overlaps more extensively with the canonical reference pattern, which contains both phrases, than does the narrower (6,7) halo. Because this effect arises from input-projection geometry rather than stored NIT memory, it provides an internal control for the footprint interpretation.

The halo asymmetry also predicts that footprint effects should extend to untrained transitions when the wider phrase's cascade contributes to the recorded NTZ pattern. This analysis comes closest to separating footprint from memory because the halo is cascade-driven and does not require NIT memory. The prediction was supported at 7→1, where the effect was measured precisely. It could not be evaluated at 0→6 because measurement variance was too large. Related neuroimaging studies have reported altered activation for semantically related but untrained material after focal lesions in language cortex (Hickok & Poeppel, 2007; Patterson et al., 2007; Saur et al., 2008), often described as spreading deficits or graded category-specific impairment.

4. Discussion

The results refine the classical cognitive-reserve hypothesis (Stern, 2002, 2009, 2012; Stern et al., 2018, 2020) within the Hexcortex model. First, IAZ fingerprint width determines how much lesion resilience can be gained from additional training. Wider fingerprints occupy larger NTZ footprints, accommodate more NIT memories before saturation, and retain count-dependent effects over a broader range. Narrower fingerprints reach saturation earlier and become count-independent above that point. The footprint × count interaction is the statistical evidence for this dependence. Second, footprint was the main determinant of focal-lesion resilience for the transitions measured with adequate precision, and its effect extended to an untrained transition through the cascade halo. These findings follow from an architecture in which memories occupy identifiable spatial locations and their distribution is constrained by upstream projection geometry. Projection geometry in biological cortex is established early and constrains later learning (Hickok & Poeppel, 2007; Seghier, 2013), making the model a source of testable hypotheses. The present study does not establish that an analogous variable determines reserve in biological cortex.

These results suggest that the classical account of clinical cognitive reserve may be underspecified. The usual question—whether more accumulated learning increases resilience—treats learning as a proxy for memory count. In the model, training affects resilience through the spatial substrate occupied by the trace and the density of memories within it. This interpretation is consistent with two areas of prior work. The maintenance/reserve/compensation framework (Cabeza et al., 2018) distinguishes accumulated cognitive content from the structural substrate that supports it. The present results identify fingerprint-footprint width as one spatial property that mediates this relationship in Hexcortex. Computational neuropsychology has also shown that lesion-symptom dissociations can arise from distributed representations whose properties cannot be reduced to memory count (Patterson et al., 2007; Plaut, 1995). Hexcortex provides an explicit substrate in which that spatial property can be measured directly.

The findings relate to three established lines of work. First, connectionist neuropsychology has long shown that distributed representations degrade more gradually under focal damage than localized representations and that recoverability depends on spatial distribution (Hinton et al., 1986; Plaut, 1995), extending the earlier mass-action principle (Lashley, 1929). The contribution here is to experimentally separate representational extent from memory quantity. In earlier demonstrations, the two often covary because additional training both adds traces and alters their distribution. The matched-count construction holds count fixed while footprint varies, and footprint fixed while count varies, linking the dissociation to a measurable substrate variable. Second, the footprint-saturation account can be placed within the cognitive-reserve taxonomy of neural reserve and neural compensation (Stern, 2002, 2009, 2012; Stern et al., 2018, 2020). It defines neural reserve in terms of the spatial spread of the representational substrate rather than general network efficiency or the quantity of stored content. This interpretation is compatible with the taxonomy but differs from accounts that define substrate-level reserve as redundancy in the number of structural routes between regions (Yoo et al., 2015). Third, the saturation effect is related to the capacity limit of associative memory (Amit et al., 1985; Hopfield, 1982) and to catastrophic interference in distributed learning systems (French, 1999; McCloskey & Cohen, 1989). The model adds a spatial constraint: the location and width of the saturating substrate are set by upstream projection geometry rather than by the learning process itself. The novel contribution is the controlled dissociation of these factors within one architecture.

Two potential clinical implications follow if the model's predictions generalize. First, reserve effects in epidemiological studies may depend on the breadth of representational substrate engaged by learning, rather than on the number of learning events alone. Extensive learning that engages widely distributed cortical territories—for example, in musicians (Draganski et al., 2004; Karni et al., 1998), bilinguals, surgeons, or people with diverse occupational histories—may provide greater resilience than equally intensive but more spatially restricted learning, even when accumulated memory counts are similar. Second, the benefit of additional training may depend on substrate-specific saturation thresholds. Above saturation, more training can add traces without increasing resilience. Within Hexcortex, this prediction can be tested using vocabularies with engineered fingerprint widths. Clinical tests could use covariate-aware measures of reserve (Stern et al., 2018, 2020). Educational attainment may be a weak proxy for representational breadth in educationally homogeneous samples, whereas tasks engaging spatially varied representations may reveal differences within otherwise similar populations (Roe et al., 2007; Soldan et al., 2017).

Four limitations constrain the interpretation. First, the vocabulary contains only two phrases, so footprint width and phrase identity vary together. Any difference attributed to the 620-cell and 880-cell footprints could instead reflect another property of the (6,7) and (1,0) representations, including fingerprint position within the input zone or overlap with the canonical reference. The footprint interpretation is supported by three converging observations: the memory-formation rate scales with fingerprint size, the count effect appears only in the wider-footprint family, as predicted by saturation, and the untrained cascade halo follows the footprint width. None of these observations replaces an independent manipulation. A decisive test will require parametric control of fingerprint width while holding phrase identity fixed, which the architecture permits. Second, the substrate uses only two phrases composed of four digits. Biological cognitive reserve reflects learning across much larger vocabularies, and saturation behavior may change at scale. Third, the random and contiguous lesion masks are schematic compared with the heterogeneous degeneration patterns of neurological disease, where vulnerability varies with cytoarchitecture and connectivity (Catani & ffytche, 2005; Mesulam, 2000; Stam, 2014). Fourth, the eight seeds in each cell sample lesion placement within one deterministic engine rather

than biological variation across individuals. The paired analysis is appropriate because the same seeds were used across substrates, but the confidence intervals describe variability in lesion placement within a single architecture. Generalization to richer vocabularies, engineered footprint widths, and more biologically grounded lesion shapes remains to be tested.

The model also raises a question about the developmental origin of footprint width. In Hexcortex, IAZ fingerprint width is determined by the surgical mapping protocol and by upstream pSTS activity. In biological systems, an analogous property would depend on the spatial distribution of projections from sensory cortices and on developmental plasticity (Hickok & Poeppel, 2007; Mountcastle, 1997; Saur et al., 2008). The present substrate cannot determine whether reserve-relevant footprint width remains modifiable through enriched experience later in development (Draganski et al., 2004; Karni et al., 1998; Petersen et al., 1998) or is largely fixed by projection geometry. The framework could be extended to test this distinction.

Cosine similarity between NTZ AC firing patterns and a canonical reference measures how lesion-induced cascade activity differs from intact-substrate activity. It belongs to the family of representational-similarity analyses used to compare neural patterns across conditions and models (Diedrichsen & Kriegeskorte, 2017; Kriegeskorte et al., 2008). We selected it because it is scale-invariant. Focal ablation reduces the magnitude of the firing vector, and a magnitude-sensitive measure such as Euclidean distance would combine loss of overall activity with a change in activity pattern. Cosine similarity separates these components and is widely used in representational similarity analysis. However, the conclusions have not been tested for metric sensitivity. Correlation distance would also remove the mean, whereas Hamming or Jaccard measures of thresholded activity would weigh sparse co-activation differently. The large effects at the well-resolved transitions are likely to persist under reasonable alternatives, but this remains to be tested. Cosine similarity is also not a behavioral measure such as sequence-recall accuracy. The conclusions, therefore, concern the resilience of neural firing patterns. Recall performance may follow the cosine measure because activity in recall-relevant NTZ columns drives word-transition prediction in this architecture, but that relationship has not yet been demonstrated. The link from pattern similarity to preserved function remains an assumption rather than a result.

Within Hexcortex, input-fingerprint geometry determined focal-lesion resilience, and the effect of accumulated memory depended on the capacity left within that geometry. The broader research program treats projection geometry as a source of computational specification that can be derived and tested independently of parameter learning (LeCun et al., 2015; Maass et al., 2002). The present experiment supports that principle for one architecture and one clinical hypothesis. Tests with parametrically engineered footprints and with other cognitive functions are needed to determine how broadly the result applies.

Credit authorship contribution statement

Fredric Narcross: Conceptualization; Methodology; Software; Investigation; Formal analysis; Data curation; Visualization; Writing – original draft; Writing – review and editing.

Robert S. Marks: Conceptualization; Supervision; Project administration; Writing – review and editing.

Declaration of competing interests

The Hexcortex architecture described in this manuscript is the subject of pending patent applications managed by B.G. Negev Technologies and Applications Ltd. (the technology-transfer office of Ben-Gurion University of the Negev), on which F.N. and R.M. are named inventors. The authors declare no other competing financial or personal relationships that could have influenced the work reported in this paper.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During preparation of this work, the authors used Claude (Anthropic) and ChatGPT (OpenAI) to support manuscript restructuring, reference reformatting, and conversion of in-text citations between citation systems. The authors reviewed and edited all AI-assisted output and take full responsibility for the final content of the article.

Data availability

Visual digit data (UCI Optical Recognition of Handwritten Digits, CC BY 4.0) are publicly available via scikit-learn's load_digits function and from the UCI Machine Learning Repository (https://doi.org/10.24432/C50P49). AudioMNIST audio recordings (MIT License, Copyright © 2018 Sören Becker) are publicly available from https://github.com/soerenab/AudioMNIST. All recordings, cache files, training scripts, and analysis code used to produce this manuscript will be made available through Ben-Gurion University's open-data archive upon publication. The Hexcortex engine source code is subject to a patent application by B.G. Negev Technologies and Applications Ltd.; access for non-commercial replication will be granted upon reasonable request to the corresponding author.

Acknowledgements

The authors thank the Department of Biotechnology Engineering and the Biosensors Laboratory at Ben-Gurion University of the Negev for institutional support.

References:

Adolphs, R. (2016). Human lesion studies in the 21st century. Neuron, 90, 1151–1153.

Alpaydin, E., & Kaynak, C. (1998). Optical recognition of handwritten digits [Dataset]. UCI Machine Learning Repository. https://doi.org/10.24432/C50P49

Amit, D. J., Gutfreund, H., & Sompolinsky, H. (1985). Storing infinite numbers of patterns in a spin-glass model of neural networks. Physical Review Letters, 55, 1530–1533.

Bartrés-Faz, D., & Arenaza-Urquijo, E. M. (2011). Structural and functional imaging correlates of cognitive and brain reserve hypotheses in healthy and pathological aging. Brain Topography, 24, 340–357.

Bates, E., Wilson, S. M., Saygin, A. P., Dick, F., Sereno, M. I., Knight, R. T., & Dronkers, N. F. (2003). Voxel-based lesion-symptom mapping. Nature Neuroscience, 6, 448–450.

Becker, S., Vielhaben, J., Ackermann, M., Müller, K.-R., Lapuschkin, S., & Samek, W. (2024). AudioMNIST: Exploring explainable artificial intelligence for audio analysis on a simple benchmark. Journal of the Franklin Institute, 361, 418–428.

Buzsáki, G. (2002). Theta oscillations in the hippocampus. Neuron, 33, 325–340.

Buzsáki, G., & Wang, X. J. (2012). Mechanisms of gamma oscillations. Annual Review of Neuroscience, 35, 203–225.

Cabeza, R., Albert, M., Belleville, S., Craik, F. I. M., Duarte, A., Grady, C. L., Lindenberger, U., Nyberg, L., Park, D. C., Reuter-Lorenz, P. A., Rugg, M. D., Steffener, J., & Rajah, M. N. (2018). Maintenance, reserve and compensation: the cognitive neuroscience of healthy aging. Nature Reviews Neuroscience, 19, 701–710.

Catani, M., Allin, M. P. G., Husain, M., Pugliese, L., Mesulam, M. M., Murray, R. M., & Jones, D. K. (2007). Symmetries in human brain language pathways correlate with verbal recall. Proceedings of the National Academy of Sciences USA, 104, 17163–17168.

Catani, M., & ffytche, D. H. (2005). The rises and falls of disconnection syndromes. Brain, 128, 2224–2239.

Colgin, L. L. (2013). Mechanisms and functions of theta rhythms. Annual Review of Neuroscience, 36, 295–312.

Damasio, H., & Damasio, A. R. (1989). Lesion analysis in neuropsychology. Oxford University Press.

Davies, M., Srinivasa, N., Lin, T.-H., Chinya, G., Cao, Y., Choday, S. H., Dimou, G., Joshi, P., Imam, N., Jain, S., Liao, Y., Lin, C.-K., Lines, A., Liu, R., Mathaikutty, D., McCoy, S., Paul, A., Tse, J., Venkataramanan, G., … Wang, H. (2018). Loihi: a neuromorphic manycore processor with on-chip learning. IEEE Micro, 38(1), 82–99.

Diedrichsen, J., & Kriegeskorte, N. (2017). Representational models: a common framework for understanding encoding, pattern-component, and representational-similarity analysis. PLOS Computational Biology, 13, e1005508.

Draganski, B., Gaser, C., Busch, V., Schuierer, G., Bogdahn, U., & May, A. (2004). Neuroplasticity: changes in grey matter induced by training. Nature, 427, 311–312.

Eichenbaum, H. (2000). A cortical-hippocampal system for declarative memory. Nature Reviews Neuroscience, 1, 41–50.

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. Cerebral Cortex, 1, 1–47.

Frankland, P. W., & Bontempi, B. (2005). The organization of recent and remote memories. Nature Reviews Neuroscience, 6, 119–130.

French, R. M. (1999). Catastrophic forgetting in connectionist networks. Trends in Cognitive Sciences, 3, 128–135.

Hickok, G., & Poeppel, D. (2007). The cortical organization of speech processing. Nature Reviews Neuroscience, 8, 393–402.

Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In D. E. Rumelhart & J. L. McClelland (Eds.), Parallel distributed processing: Explorations in the microstructure of cognition (Vol. 1, pp. 77–109). MIT Press.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. Proceedings of the National Academy of Sciences USA, 79, 2554–2558.

Karni, A., Meyer, G., Rey-Hipolito, C., Jezzard, P., Adams, M. M., Turner, R., & Ungerleider, L. G. (1998). The acquisition of skilled motor performance: fast and slow experience-driven changes in primary motor cortex. Proceedings of the National Academy of Sciences USA, 95, 861–868.

Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis: connecting the branches of systems neuroscience. Frontiers in Systems Neuroscience, 2, 4.

Lashley, K. S. (1929). Brain mechanisms and intelligence: A quantitative study of injuries to the brain. University of Chicago Press.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521, 436–444.

Lisman, J. E., & Jensen, O. (2013). The theta-gamma neural code. Neuron, 77, 1002–1016.

Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. Neural Computation, 14, 2531–2560.

Markram, H., Muller, E., Ramaswamy, S., Reimann, M. W., Abdellah, M., Sanchez, C. A., Ailamaki, A., Alonso-Nanclares, L., Antille, N., Arsever, S., Kahou, G. A., Berger, T. K., Bilgili, A., Buncic, N., Chalimourda, A., Chindemi, G., Courcol, J.-D., Delalondre, F., Delattre, V., … Schürmann, F. (2015). Reconstruction and simulation of neocortical microcircuitry. Cell, 163, 456–492.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychological Review, 102, 419–457.

McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: the sequential learning problem. Psychology of Learning and Motivation, 24, 109–165.

Mesulam, M.-M. (2000). Principles of behavioral and cognitive neurology (2nd ed.). Oxford University Press.

Mountcastle, V. B. (1997). The columnar organization of the neocortex. Brain, 120, 701–722.

Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. Nature Reviews Neuroscience, 8, 976–987.

Petersen, S. E., van Mier, H., Fiez, J. A., & Raichle, M. E. (1998). The effects of practice on the functional anatomy of task performance. Proceedings of the National Academy of Sciences USA, 95, 853–860.

Plaut, D. C. (1995). Double dissociation without modularity: evidence from connectionist neuropsychology. Journal of Clinical and Experimental Neuropsychology, 17, 291–321.

Roe, C. M., Xiong, C., Miller, J. P., & Morris, J. C. (2007). Education and Alzheimer disease without dementia: support for the cognitive reserve hypothesis. Neurology, 68, 223–228.

Rorden, C., & Karnath, H.-O. (2004). Using human brain lesions to infer function: a relic from a past era in the fMRI age? Nature Reviews Neuroscience, 5, 813–819.

Roy, K., Jaiswal, A., & Panda, P. (2019). Towards spike-based machine intelligence with neuromorphic computing. Nature, 575, 607–617.

Saur, D., Kreher, B. W., Schnell, S., Kümmerer, D., Kellmeyer, P., Vry, M.-S., Umarova, R., Musso, M., Glauche, V., Abel, S., Huber, W., Rijntjes, M., Hennig, J., & Weiller, C. (2008). Ventral and dorsal pathways for language. Proceedings of the National Academy of Sciences USA, 105, 18035–18040.

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. Journal of Neurology, Neurosurgery & Psychiatry, 20, 11–21.

Seghier, M. L. (2013). The angular gyrus: multiple functions and multiple subdivisions. The Neuroscientist, 19, 43–61.

Soldan, A., Pettigrew, C., Cai, Q., Wang, J., Wang, M.-C., Moghekar, A., Miller, M. I., Albert, M., & BIOCARD Research Team. (2017). Cognitive reserve and long-term change in cognition in aging and preclinical Alzheimer's disease. Neurobiology of Aging, 60, 164–172.

Squire, L. R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. Psychological Review, 99, 195–231.

Squire, L. R., & Wixted, J. T. (2011). The cognitive neuroscience of human memory since H.M. Annual Review of Neuroscience, 34, 259–288.

Stam, C. J. (2014). Modern network science of neurological disorders. Nature Reviews Neuroscience, 15, 683–695.

Stern, Y. (2002). What is cognitive reserve? Theory and research application of the reserve concept. Journal of the International Neuropsychological Society, 8, 448–460.

Stern, Y. (2009). Cognitive reserve. Neuropsychologia, 47, 2015–2028.

Stern, Y. (2012). Cognitive reserve in ageing and Alzheimer's disease. Lancet Neurology, 11, 1006–1012.

Stern, Y., Arenaza-Urquijo, E. M., Bartrés-Faz, D., Belleville, S., Cantilon, M., Chételat, G., Ewers, M., Franzmeier, N., Kempermann, G., Kremen, W. S., Okonkwo, O., Scarmeas, N., Soldan, A., Udeh-Momoh, C., Valenzuela, M., Vemuri, P., Vuoksimaa, E., & Reserve, Resilience and Protective Factors PIA Empirical Definitions and

Conceptual Frameworks Workgroup. (2018). Whitepaper: Defining and investigating cognitive reserve, brain reserve, and brain maintenance. Alzheimer's & Dementia. Advance online publication. https://doi.org/10.1016/j.jalz.2018.07.219

Stern, Y., Barnes, C. A., Grady, C., Jones, R. N., & Raz, N. (2020). A framework for concepts of reserve and resilience in aging. Neurobiology of Aging, 83, 124–129.

Valenzuela, M. J., & Sachdev, P. (2006). Brain reserve and dementia: a systematic review. Psychological Medicine, 36, 441–454.

Whalley, L. J., Deary, I. J., Appleton, C. L., & Starr, J. M. (2004). Cognitive reserve and the neurobiology of cognitive aging. Ageing Research Reviews, 3, 369–382.

Yoo, S. W., Han, C. E., Shin, J. S., Seo, S. W., Na, D. L., Kaiser, M., Jeong, Y., & Seong, J. K. (2015). A network flow-based analysis of cognitive reserve in normal ageing and Alzheimer's disease. Scientific Reports, 5, Article 10057. https://doi.org/10.1038/srep10057