
Sutra: Tensor-Op RNNs as a Compilation Target for Vector Symbolic Architectures

Anonymous Author(s)

Affiliation

Address

email

1

2 Abstract

3 **Sutra** is a typed, purely functional programming language whose compiled forward pass is a Py-
4 Torch neural network. The compiler beta-reduces the whole program — primitives, control flow,
5 string I/O — to one fused tensor-op graph over a frozen embedding substrate. Rotation binding,
6 unbind, bundle, polynomial Kleene three-valued logic, and tail-recursive loops all lower to tensor
7 operations; the Kleene connectives are Lagrange-interpolated polynomials exact on the $\{-1, 0, +1\}$
8 truth grid.

9 Validation is one fact tested two ways. (1) The same program runs on four frozen embeddings span-
10 ning two modalities — three text encoders (nomic-embed-text, all-minilm, mxbai-embed-large) and
11 one protein language model (ESM-2) — and decodes bundles at 100% accuracy through width $k=8$
12 on every substrate, where the textbook Hadamard product has already collapsed (2.5% on mxbai-
13 embed-large, 7.5% on all-minilm). (2) PyTorch autograd flows through the actually compiled graph:
14 a fuzzy-rule classifier written in .su trains from random init ($18.7 \pm 9.5\%$; chance = 20%, five
15 classes) to $100.0 \pm 0.0\%$ (three seeds) by backpropagating through the emitted graph, the symbolic
16 source unmodified. A weighted variant additionally trains a scalar cosine gain and writes it back
17 into the .su source as a numeric literal; recompiling reproduces the trained behaviour to $\approx 2 \times 10^{-7}$
18 per logit, so the trained model is itself legible, recompilable code.

19 The same artifact is therefore both a logic program and a trainable neural network.

20

21 1 Introduction

22 A frozen embedding model maps strings (or amino-acid sequences, or any other input the model
23 was trained on) into a deterministic continuous vector space. Given such a substrate, two technical
24 questions follow:

- 25 1. **Which operations on these embeddings are reliable enough to be used as primitives of**
26 a compositional algebra over the substrate's vector space?
- 27 2. **What is the correct binding operation?** Hyperdimensional computing's textbook bind
28 operators (Hadamard product, circular convolution) were derived assuming hypervectors
29 drawn from a controlled random distribution. Frozen LLM embeddings are not such a dis-
30 tribution: they are strongly *anisotropic*, concentrating in a narrow cone so cosine similarity
31 is compressed and inflated even between unrelated items (Ethayarajh 2019). §3.2 measures
32 four substrates and reports that rotation binding decodes at 100% accuracy through bundle
33 widths where Hadamard has already collapsed.

34 This paper answers both questions in the form of a working programming language, **Sutra**, whose
 35 primitives are these consolidated operations and whose compiled forward pass is a PyTorch neural
 36 network. The naming: **Sutra** is the Sanskrit *sūtra* (thread, rule, aphorism), the term for Pāṇini's
 37 foundational Sanskrit grammar.

38 1.1 Contributions

39 The four core technical contributions of this paper are:

40 1. **Polynomial fuzzy logic via Lagrange interpolation of Kleene's three-valued truth ta-**
 41 **bles.** The truth axis encodes $T = +1, U = 0, F = -1$. On the discrete $\{-1, 0, +1\}$
 42 grid, the Kleene connectives are AND = min, OR = max, NOT = $-\cdot$. The min/max
 43 forms (the standard Gödel t-norm/t-conorm choice; Hájek 1998) are non-differentiable at
 44 the diagonal $a = b$, which breaks gradient flow when connectives compose with the tensor-
 45 op graph (van Krieken, Acar & van Harmelen 2022 survey the issue across t-norm-derived
 46 neural-symbolic operators). Sutra resolves this by Lagrange-interpolating each connective
 47 as a polynomial that is exact on the 3×3 Kleene grid and C^∞ elsewhere:

$$\begin{aligned} \text{AND}(a, b) &= \frac{1}{2}(a + b + ab - a^2 - b^2 + a^2b^2) \\ \text{NAND}(a, b) &= \frac{1}{2}(-a - b - ab + a^2 + b^2 - a^2b^2) \\ \text{OR}(a, b) &= \frac{1}{2}(a + b - ab + a^2 + b^2 - a^2b^2) \\ \text{NOR}(a, b) &= \frac{1}{2}(-a - b + ab - a^2 - b^2 + a^2b^2) \\ \text{NOT}(a) &= -a \\ \text{XOR}(a, b) &= -ab \\ \text{XNOR}(a, b) &= ab \end{aligned}$$

48 {AND, OR, NOT} is functionally complete for the Kleene fragment; NAND and NOR
 49 are just $-\text{AND}$ and $-\text{OR}$ (negation is $-\cdot$), and XOR/XNOR collapse to a single multi-
 50 plicative term because their interpolant is zero whenever either input is U and bilinear in
 51 the $\{-1, +1\}$ corners. Every Kleene-valid connective is therefore a polynomial tensor-op-
 52 graph fragment, gradient-compatible, branchless, and exact on the discrete-logic regime.
 53 A symbolic if-then rule built from these gates is one fused subgraph that PyTorch autograd
 54 backprops through end-to-end (§3.6).

55 2. **Beta reduction to a substrate-pure tensor-op graph.** The compiler inlines stdlib operator
 56 definitions, beta-reduces through bound names, then runs an algebraic-simplification pass
 57 over the residual. What's left is a fused tensor-op graph (matmul / element-wise / nonlinear)
 58 with no named bindings or function calls. Three concrete moves go beyond standard inlin-
 59 ing + constant folding: conditionals lower to soft-mux polynomials ($\frac{1+\text{cond}}{2}a + \frac{1-\text{cond}}{2}b$)
 60 so the compiled artifact has no if opcodes; Haar-orthogonal binding rotations R_{role} are
 61 materialized at compile time so runtime bind is one matmul against a constant matrix;
 62 canonical synthetic axes are assigned compile-time so every primitive-type read/write is a
 63 known index, not a hashtable lookup. §4.2 traces this lowering stage-by-stage on a concrete
 64 program; the compilation pipeline as a whole is diagrammed in Figure~3.

65 3. **Tail recursion as the loop primitive.** Loops are tail-recursive function declarations
 66 (do_while, while_loop, iterative_loop, foreach_loop) whose body's return
 67 NAME(args) becomes the recurrent step. Each loop compiles to a soft-halt RNN cell with
 68 substrate-pure halt detection (heaviside \rightarrow cumulative monotone halt \rightarrow soft-mux state
 69 freeze). The body of every loop tick is one straight-line tensor pipeline with no in-graph
 70 branches; a thin Python while True: ... break driver wraps the body and terminates
 71 when the halt scalar saturates (§3.4). The state vector is fixed-width across iterations, **O(1)**
 72 state, **O(N)** compute, **O(N)** gradient tape during training, where N is iterations actually
 73 executed.

74 4. **Synthetic-dimension rotation binding as an angular hash map.** The compiler reserves
 75 a synthetic block of canonical dimensions and uses Haar-orthogonal rotations seeded from
 76 the role's content hash to bind keys to slots. We are unaware of any prior use of a high-
 77 dimensional rotation pattern as the substrate for a functional hash-map primitive.

78 These four primitives integrate into a single working compiler that lowers `.su` source to a self-
79 contained PyTorch module on CPU or CUDA. Program inputs and outputs are embeddings in the
80 substrate’s vector space; a compile-time codebook (implemented with an embedded vector database,
81 §3.5) handles the convenience of source-level string literals and nearest-string output.

82 1.2 The substrate is the architecture target

83 A Sutra program is compiled for an *embedding-space architecture*, the way a C program is compiled
84 for x86 and a CUDA kernel for an NVIDIA SM. The embedding model fixes dimensionality, the
85 geometry of the semantic block, and the meaning of every basis-vector lookup; swap the model and
86 the same source recompiles to a different `.sdb` codebook against a different geometry. The substrate
87 need not be an LLM, it can be any network producing a dense vector representation, including
88 the hidden state of a trained model. §3.2’s ESM-2 protein-LM row demonstrates this substrate-
89 agnostically.

90

91 2 Related Work

92 2.1 Vector Symbolic Architectures

93 VSA is a family of algebraic frameworks for computing with high- dimensional vectors (Kanerva
94 2009; Plate 1995; Gayler 2003). The standard VSA development assumes hypervectors drawn from
95 a controlled random distribution designed for the algebra; bind is typically Hadamard product or
96 circular convolution. Frozen LLM embedding spaces are not designed for VSA: they are anisotropic
97 (Ethayarajh 2019; Gao et al. 2019; Mu et al. 2018) — representations concentrate in a narrow
98 cone, so cosine similarity has low dynamic range and the textbook bind operations do not transfer
99 cleanly. The same anisotropy makes an unweighted cosine read a weak rule signal, which is why
100 §3.6 trains the embeddings the rules compare against rather than relying on raw cosine. Rotation
101 binding (`R_role @ filler` for a role-seeded Haar-random orthogonal `R_role`) is the choice that
102 worked across the substrates we tested, and is what Sutra uses today; §3.2 reports the per-substrate
103 measurements supporting that choice.

104 The closest software peer in the VSA space is **TorchHD** (Heddes et al. 2023), a PyTorch library that
105 exposes VSA primitives (`bind`, `bundle`, `similarity`) as tensor operations. Sutra and TorchHD differ
106 on what the user writes and what the compiler does:

- 107 • **TorchHD is a *library*.** The user writes Python code that calls TorchHD primitives; control
108 flow is host-side Python; there is no source-language layer above the primitives, no compile
109 step, and no algebraic reduction across primitive calls. Each primitive call is a tensor op,
110 but the program itself is a Python function with whatever control flow the user wrote.
- 111 • **Sutra is a *language with a compiler*.** The user writes `.su` source which the compiler
112 beta-reduces to a substrate-pure tensor-op graph (§1.1-2): a single straight-line graph of
113 `matmul` / `element-wise` / `nonlinear` ops with no Python control flow. Loops are tail-recursive
114 function declarations that lower to soft-halt RNN cells; conditionals are differentiable fuzzy
115 interpolations rather than Python `if`. Hash-map structure is implemented via synthetic-
116 dimension rotation, not via a host-side dictionary.

117 A second axis where Sutra differs from existing HDC software is **string I/O**. TorchHD and similar
118 libraries expose the algebra over user-supplied hypervectors; the user maintains a `dict[str, hypervector]`
119 and an explicit codebook tensor by hand. Sutra’s compile-time codebook (§3.5) closes that loop:
120 every embedded string in `.su` source is embedded once at compile time via the configured frozen LLM,
121 stored in the project’s `.sdb` codebook, and decoded at the program output
122 via `nearest_string`. The frozen-LLM embedding is load-bearing, random hypervectors yield a
123 working VSA algebra with no I/O story.

124 The structural differences (Sutra contains no Python, the string-to-vector map and codebook are
125 constructed by the compiler rather than by the user, and the whole program reduces to a single fused
126 tensor-op graph) are differences in artifact shape, not library speed.

127 2.2 Comparison to other neuro-symbolic languages

128 The closest neuro-symbolic-language peers are **Scallop** (Li et al. 2023, Datalog with provenance-
129 semiring differentiability), **DeepProbLog** (Manhaeve et al. 2018, ProbLog with neural predicates),
130 **Logic Tensor Networks** (Badreddine et al. 2022, first-order logic compiled to t-norm losses), and
131 **NeurASP** (Yang et al. 2020, Answer Set Programming with neural predicates). All share a two-
132 stage perception-then-reasoning shape: a neural model extracts discrete symbols from raw input,
133 and a symbolic program reasons over those symbols. Sutra's shape is different at this architectural
134 level: the substrate is a continuous embedding space throughout, primitives operate on vectors end-
135 to-end, and the whole program (including what would be the logic program in Scallop) compiles
136 to a single fused tensor-op graph through beta reduction. There is no discrete symbolic stratum to
137 extract into or reason over; differentiability is inherited from the tensor-op graph itself, not from
138 a provenance annotation on a relational query. The two are good at different problem structures:
139 Scallop and its peers when the problem is naturally relational and perception cleanly factors out;
140 Sutra when computation is best expressed as algebra on vectors over a substrate the program reads
141 strings into and decodes strings out of.

142 The closest HDC peer with compiler infrastructure is **HDCC** (Vergés et al. 2023), a description-file
143 DSL targeting self-contained C for embedded classification, random/level hypervectors only, no gen-
144 eral control flow, scoped to classification. **TorchHD** and **OpenHD / HD Torch** are libraries without a
145 language-level loop primitive. To the authors' knowledge (literature reviewed through early 2026),
146 no published HDC system combines (a) one fused tensor-op graph as compile target, (b) HDC prim-
147 itives as the operations, (c) a frozen externally-trained vector embedding space as the substrate, and
148 (d) tail-recursive loops compiled to soft-halt RNN cells with constant state-vector width in recursion
149 depth. The combination is what distinguishes Sutra, not any one of those properties in isolation.

150 2.3 Fuzzy logic and neuro-fuzzy systems

151 Sutra's polynomial connectives sit downstream of fuzzy set theory (Zadeh 1965) and the neuro-fuzzy
152 lineage that makes fuzzy inference trainable — adaptive-network fuzzy inference systems (Jang
153 1993) and the broader fuzzy-neural-network family (Buckley & Hayashi 1994). Those systems
154 *learn* the fuzzy logic itself: membership functions and rule parameters are the trainable object. Sutra
155 inverts this. Its connectives are fixed Lagrange interpolants of Kleene's three-valued tables (§1.1-1),
156 exact on the discrete grid and never tuned; the only learnable parameters are the embeddings the
157 frozen rule graph evaluates against (§3.6). Sutra is therefore closer in spirit to the differentiable t-
158 norm analysis of van Krieken, Acar & van Harmelen (2022) than to membership-function learning,
159 and differs from both by compiling the connectives into a single substrate-pure tensor-op graph
160 rather than evaluating them in a host interpreter.

161 2.4 Differentiable Programming, AOT Compilation, and Knowledge

162 Compilation

163 The closest design ancestors are partial-evaluation systems that specialize programs at compile time
164 (the Futamura projections), differentiable programming systems that treat programs as differentiable
165 functions (JAX), AOT compilation of neural networks (TVM, XLA), and knowledge compilation in
166 symbolic AI (Darwiche & Marquis 2002). Sutra differs from each: TVM/XLA start from a network,
167 not toward one; JAX treats programs as differentiable but does not bake source literals into weights;
168 partial evaluation specializes for compile-time-known values but does not target a neural-network-
169 shaped artifact; knowledge compilation targets Boolean circuits, not continuous embedding spaces.
170 Sutra's combination (fold source literals into the weight structure, compile control flow to RNN cells,
171 run the whole program as one tensor-op graph over a *continuous* substrate) is the novel position.

172

173 3 Consolidation into Canonical Primitives

174 The central design move: hold the operation interface fixed and pick a binding implementation
175 that works on dense externally-trained substrates. Standard VSA's Hadamard product fails here,

176 elementwise multiplication of correlated real-valued vectors produces destructive crosstalk on bun-
 177 dled retrieval (§3.2 measures this directly). Rotation binding works: each role gets a Haar-random
 178 orthogonal R_{role} seeded by $\text{hash}(\text{role})$, and $\text{bind}(\text{role}, \text{filler}) = R_{\text{role}} @ \text{filler}$ is
 179 invertible (unbind is the transpose) and well-conditioned. The compiler caches R_{role} per-role at
 180 module init so runtime bind is a single matmul against a precomputed matrix.

181 Each subsection below is tagged *method* (the consolidated primitives, extended-state-vector layout,
 182 loop cell, codebook, and type surface) or *experiment* (the cross-substrate capacity measurements,
 183 §3.2 / §3.2.1, and the end-to-end differentiable-training result, §3.6). The two are interleaved by
 184 topic but labelled so the evaluation is separable from the design.

185 3.1 Notation — method

186 We work in \mathbb{R}^d with d the substrate's embedding dimension (768 for nomic-embed-text). Every
 187 value has the layout [semantic | synthetic]. The seven primitive operations: $\text{bind}(r, f) = R_r f$
 188 where $R_r = \text{QR}(\text{hash}(r))$, Q is Haar-orthogonal, $\text{unbind}(r, v) = R_r^T v$, $\text{bundle}(x, y) = (x +$
 189 $y) / (\|x + y\| + \varepsilon)$, $\text{similarity}(x, y) = (x \cdot y) / (\|x\| \|y\| + \varepsilon)$, $\text{normalize}(v) = v / (\|v\| + \varepsilon)$, the
 190 Lagrange Kleene gates as in §1.1-1, and the soft-halt cell of §3.4. Full signature/definition table and
 191 the soft-halt cell update equations are in Appendix A.

192 3.2 Capacity of rotation versus Hadamard binding across substrates — experiment

193 We measure decode accuracy as a function of bundle width k on real embeddings across
 194 four substrates spanning two modalities: three frozen LLM text encoders (nomic-embed-
 195 text, all-minilm, mxbai-embed-large) and one frozen protein language model (ESM-2 small,
 196 facebook/esm2_t6_8M_UR50D). LLM substrates embed an 84-word noun vocabulary; the ESM-2
 197 substrate embeds an 84-sequence amino-acid vocabulary (full protocol in Appendix C). For each
 198 bundle width and binding scheme we run 10 trials, sampling k random (role, filler) pairs without re-
 199 placement, forming the bundle, and decoding by $\text{unbind} + \text{argmax-cosine}$ against the full codebook.
 200 *Rotation binding* uses a role-seeded Haar-orthogonal R_{role} ; *Hadamard binding* is the textbook
 201 elementwise product (MAP-VSA).

202 Cross-substrate decode accuracy at representative widths (full $k \in \{2, 4, 8, 16, 24, 32, 48\}$ sweeps
 203 in Appendix C):

substrate (dim)	rotation k=8	rotation k=48	Hadamard k=8	Hadamard k=48
nomic-embed-text (768)	100.0%	93.3%	87.5%	48.3%
all-minilm (384)	100.0%	42.3%	7.5%	1.7%
mxbai-embed-large (1024)	100.0%	72.1%	2.5%	1.0%
ESM-2 (320)	100.0%	44.2%	28.7%	4.2%

204 ESM-2 (Lin et al., Science 2023) is a protein language model trained on UniRef with
 205 no natural-language exposure; the same rotation-vs-Hadamard pattern reproduces in that
 206 modality. Rotation reversibility round-trip across all four substrates: mean $\|\text{unbind}(R,$
 207 $\text{bind}(R, x)) - x\| = 1.5 \times 10^{-15}$ (floating-point round-off, Q orthogonal). Reproduction:
 208 `experiments/rotation_binding_capacity_{llm,bioinformatics}.py`.

209 3.2.1 Noise accumulation across chained bind/unbind cycles — experiment

210 The §3.2 protocol measures one bind+bundle+unbind cycle. Nested records (a recovered filler be-
 211 coming the role of a sub-record) add bundle noise per level. We measured this directly: chain
 212 lengths $L \in \{1, 2, 4, 8, \dots\}$, 20 trials, bundle width 4. Raw accuracy holds at 100% through $L=2$ on
 213 every substrate and falls to chance (1/84) by $L=8$. The demonstrated regime is therefore single-cycle
 214 records, which matches the shape of the `role_filler_record`, `knowledge_graph`, and `predicate-`
 215 `lookup` demos. Pure rotation chains without per-step distractor bundling remain exact (round-trip
 216 1.5×10^{-15} per cycle), so the noise mechanism here does not apply to the soft-halt loop cell of §3.4.
 217 Reproduction script: `experiments/crosstalk_chain.py`; full per-substrate L -sweep tables in
 218 Appendix D.

219 **3.3 The extended-state-vector layout — method**

220 Every value carries a fixed [semantic | synthetic] layout: the d-dimensional semantic block
 221 holds the substrate embedding for vector-shaped values, and a small synthetic block reserves canon-
 222 ical axes for primitive types (real, imag, truth, char) and a loop-completion flag, with the remaining
 223 axes paired into 2D Givens planes for variable slots. Default at $d = 768$ (nomic-embed-text): a 100-
 224 dim synthetic block accommodates the five canonical axes plus 47 disjoint slots. Rotation binding
 225 is block-diagonal across the split (Q_role is Haar-random in the semantic block, identity on the
 226 synthetic block), so the synthetic axes pass through bind/unbind unchanged, a fuzzy-truth scalar can
 227 coexist with a semantic vector inside the same value without bind smearing them. Full per-axis
 228 purpose table and slot allocator details in Appendix B.

229 **3.4 First-class loops as RNN cells — method**

230 Runtime data-dependent loops compile to **self-halting RNN cells**. Each tick: snapshot pre-step
 231 state, evaluate halt on the substrate (truth-axis read \rightarrow heaviside \rightarrow cumulative saturating sum to
 232 halted), run the cell body, soft-mux between pre- and new-step state by halted. A Python while
 233 True: driver breaks the moment halted saturates; this is the only host-side branch in the loop
 234 machinery. Inside the cell body, every operation is a substrate tensor op. No compile-time iteration
 235 cap, programs terminate when their halt condition fires. Standard PyTorch tracing handles a Python
 236 while-loop wrapping pure tensor ops; autograd records each iteration as it executes, which is the
 237 mechanism §3.6 relies on for backprop through the cell. Figure~1 visualizes one tick.

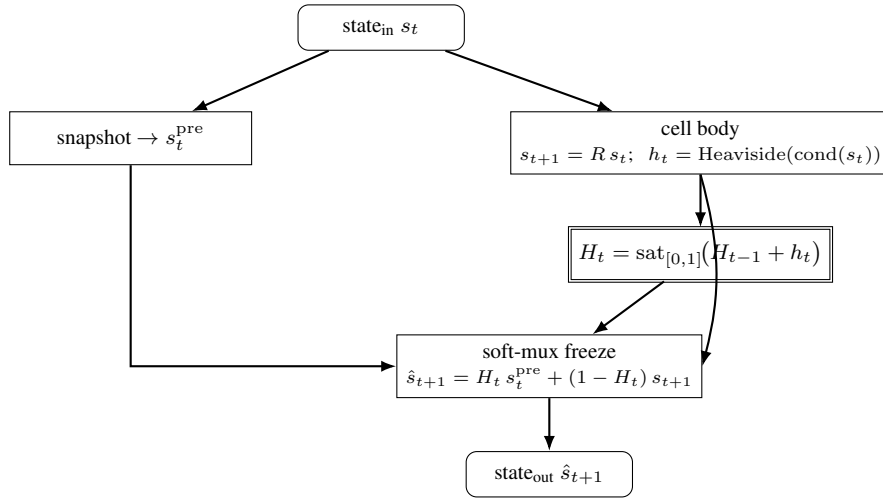


Figure 1: Per-tick dataflow of the soft-halt RNN cell. Once H_t saturates at 1, the soft-mux output equals s_t^{pre} , the loop has frozen. The cumulative halt H_t acts as a boundary read of the same shape as the codebook decode (§3.5).

238 **Constant memory in recursion depth.** The state vector is fixed-width and shared across iterations,
 239 so a tail-recursive loop consumes $O(1)$ memory in the state vector regardless of trip count. Com-
 240 pute is $O(N)$ and the autograd tape during training is $O(N)$ in iterations actually executed (standard
 241 PyTorch, freed after backward). To the authors' knowledge no other HDC system or compiler ex-
 242 poses user-program-level recursion: HDCC is scoped to classification pipelines, TorchHD requires
 243 the user to write Python loops over hypervectors. The recurrent shape that emerges is the rational-
 244 weight RNN form whose computational power Siegelmann & Sontag (1992) characterized.

245 **3.5 I/O is in the embedding space; the codebook is a comfort layer — method**

246 A Sutra program's inputs and outputs are embeddings in the substrate's vector space. Strings are a
 247 convenience for writing source-level literals: every string literal in .su source is embedded once
 248 at compile time and stored in a **codebook** (implemented as an embedded vector database with an
 249 HNSW index, on disk as a .sdb file shipped alongside the compiled module). At the program's
 250 output boundary, the runtime `decode_VSA.nearest_string(query)` maps a query embedding to

251 the nearest stored string when the program's caller wants a string back. Calling the codebook at
 252 this boundary is shape-equivalent to calling PyTorch for a matmul, neither is the kind of host-side
 253 control flow substrate purity forbids. Implementation details (RDF triple layout, HNSW parameters,
 254 .sdb file format, complexity analysis) are in Appendix E.

255 3.6 End-to-end differentiable training through Sutra operations — experiment

256 Because every Sutra primitive compiles to a differentiable tensor operation, the compiled graph
 257 supports standard PyTorch `loss.backward()` without modification. We verify this by training
 258 learnable parameters through a fuzzy-logic classifier built entirely from Sutra operations.

259 **Setup.** The classifier is written in .su and compiled by the PyTorch codegen; the emitted module's
 260 rule function is the compiler output — `_VSA.similarity` composed with the Lagrange–Kleene
 261 AND/NOT polynomials, with no hand-written reimplement. Five semantic classes, ten words
 262 each (50 inputs), embedded via nomic-embed-text (768-d, frozen). Five learnable prototype vectors
 263 are initialized randomly. The program compiled at $K = 5$ (generated by `gen_rule_su`; line-
 264 wrapped here, .su is whitespace-insensitive) is:

```
265 function fuzzy rule(vector x, vector own,
266                   vector o0, vector o1, vector o2, vector o3) {
267     return similarity(x, own)
268         && !similarity(x, o0) && !similarity(x, o1)
269         && !similarity(x, o2) && !similarity(x, o3);
270 }
```

271 so each class score is the compiled fuzzy rule

$$\text{rule}_i = \text{AND}\left(\text{sim}(x, p_i), \bigwedge_{j \neq i} \text{NOT}(\text{sim}(x, p_j))\right)$$

272 with the AND-of-NOTs left-folded across the $K - 1$ other classes (at $K = 5$, four NOT terms
 273 folded under one AND). Full-batch cross-entropy over the five compiled rule scores drives Adam
 274 (Kingma & Ba 2015; defaults $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\varepsilon = 10^{-8}$, learning rate 0.01) on the prototype
 275 embeddings, backpropagating through the emitted graph. 30 epochs, three seeds (0–2).

276 **Results (3 seeds).** From random-init accuracy at chance ($18.7 \pm 9.5\%$; chance = 20%), training
 277 reaches $100.0 \pm 0.0\%$ (mean \pm s.d. over seeds 0–2; every seed reaches 100%); cross-entropy loss
 278 falls to ≈ 0.43 . Every prototype receives a nonzero gradient, verified to propagate through the
 279 emitted graph (`_VSA.similarity` \rightarrow the emitted Lagrange–Kleene NOT/AND \rightarrow cross-entropy),
 280 not through a reimplement.

Phase	Accuracy (mean \pm s.d., $n=3$)
Before (random)	$18.7 \pm 9.5 \%$
After (30 ep)	$100.0 \pm 0.0 \%$

281 This experiment isolates gradient flow through the *compiled* symbolic graph: it trains and evaluates
 282 on the same 30-word set and reports in-sample accuracy purely as verification that backprop reaches
 283 every learnable prototype through the compiler's emitted ops — not a generalization claim; no held-
 284 out split. Only before/after accuracy was logged for the compiled run; per-epoch data was not
 285 recorded, so no intermediate-epoch curve is plotted.

286 Figure~2 shows the rule pipeline at $K = 3$, the smallest instance; the experiment uses the same
 287 shape at $K = 5$ — five cosine-similarity nodes against five learnable prototypes, the AND-of-NOTs
 288 folded over the four others. Each `sim_i` enters one branch of the AND-tree (rule i takes `sim_i`
 289 directly and `NOT(sim_j)` for $j \neq i$); the rule scores are stacked, temperature-scaled, softmaxed,
 290 and cross-entropied. Every node is a compiler-emitted tensor op — no Python branches, no string-
 291 keyed lookup — so backprop reaches every prototype through *the same compiled graph that runs at*
 292 *inference* (literally so: the graph is the codegen output, not a reimplement).

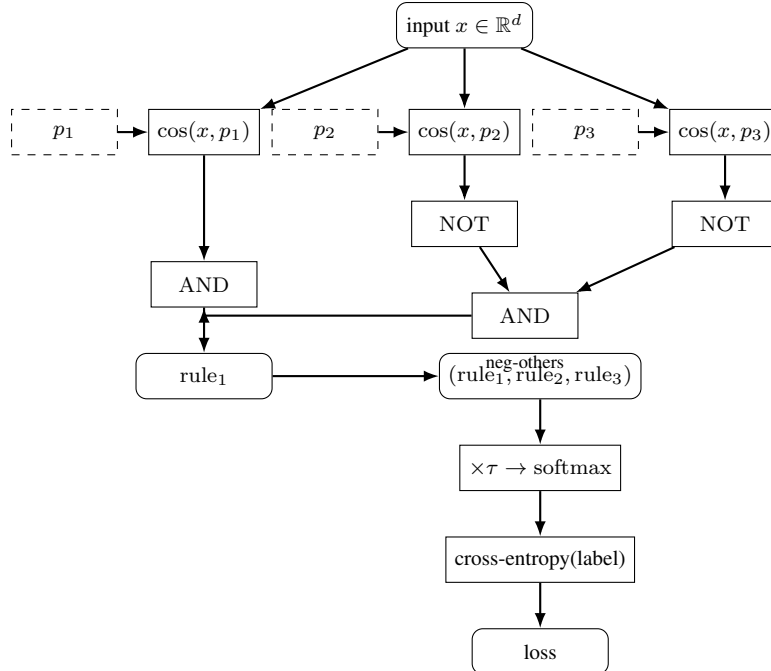


Figure 2: The $K = 3$ rule pipeline. Solid boxes are PyTorch tensor ops; dashed boxes are learnable prototypes. The AND in the leftmost branch combines $\cos(x, p_1)$ with the AND-of-NOTs over the other classes; rule₂ and rule₃ (omitted for clarity) have the symmetric shape. Every edge is a tensor; backprop reaches each p_i through this graph.

293 The AND-of-NOTs chain is a tensor pipeline that could naively saturate or vanish gradients; empiri-
 294 cally it does not — every prototype receives a nonzero gradient and the five classes separate perfectly
 295 within 30 epochs, the symbolic program text unchanged across training. Standard `torch.autograd`
 296 suffices (no Sutra-specific autograd machinery) because the compiler emits only operations PyTorch
 297 already differentiates. The `.su` compiles once; the emitted `rule` is then evaluated over the batch
 298 with `torch.vmap` — a transform that runs the *same* compiled ops with a batch axis, not a reim-
 299 plementation. Before training, the harness asserts the batched and per-sample evaluations agree
 300 within 10^{-4} on identical inputs and parameters, so the speedup is provably the identical com-
 301 piled computation. Under this path the 5-class / 50-word / 30-epoch / 3-seed run takes ≈ 230 s
 302 on CPU and yields the numbers above ($18.7 \pm 9.5\% \rightarrow 100.0 \pm 0.0\%$, three seeds). The earlier
 303 per-sample driver — one emitted `rule` call per class per input in a Python loop — produced the
 304 bit-identical result but took ≈ 6.2 h; that cost was Python-level call and per-sample autograd-graph
 305 overhead, not the compiled tensor math, and does not bound the achievable scale. Reproduction:
 306 `experiments/differentiable_training_compiled.py --batched` (compiles the `.su` once,
 307 `vmaps` the emitted `rule`, backprops the prototypes through the emitted graph; the equivalence as-
 308 sertion is on by default).

309 3.7 Trained weights compile back to legible source — experiment

310 Frozen LLM embedding spaces are *anisotropic*: learned token representations occupy a narrow cone
 311 rather than filling the sphere (the representation-degeneration effect; Gao et al. 2019, Ethayarajh
 312 2019), so raw cosine similarities are squeezed into a compressed positive band — even unrelated
 313 pairs score well above zero. The Kleene AND/NOT decision in §3.6 depends on the *gap* between the
 314 matching cosine and the off-class cosines; when anisotropy compresses that band, the polynomial
 315 connectives have little dynamic range to separate classes. This motivates a weighted similarity —
 316 `Equals(A, B, w)` — in which a single learned scalar gain w rescales the cosine term *before* the
 317 Kleene polynomials act, directly counteracting the anisotropic compression:

$$\text{rule}_i = \text{AND}\left(w \cdot \text{sim}(x, p_i), \bigwedge_{j \neq i} \text{NOT}(w \cdot \text{sim}(x, p_j))\right)$$

318 **Setup.** The rule is written in `.su` with `w` declared as a number parameter and compiled by the
 319 PyTorch codegen; the emitted code uses `w` as a plain scalar multiplying the compiler-emitted
 320 similarity, so `w` both (i) trains through the emitted graph and (ii) is a single scalar that can
 321 be written back into source as a literal. Three semantic classes, eight words each (24 inputs), nomic-
 322 embed-text (768-d, frozen). The scalar gain `w` (initialized 1.0) **and** the three prototype vectors are
 323 trained through the emitted graph; full-batch cross-entropy, Adam (lr 0.02), 30 epochs, two seeds
 324 (0–1).

325 **Results (2 seeds).** From random-init accuracy at chance ($33.3 \pm 5.9\%$; chance = 33.3%), training
 326 reaches $100.0 \pm 0.0\%$ (every seed). The learned gain converges to $w^* = 1.434 \pm 0.004$ — both
 327 seeds move it the same way, from 1.0 to ≈ 1.43 , sharpening the compressed cosine band rather than
 328 leaving it at the identity (a learned rescaling, not a tautology).

Phase	Accuracy ($n=2$)	Learned gain w^*
Before (random)	$33.3 \pm 5.9 \%$	1.000 (init)
After (30 ep)	$100.0 \pm 0.0 \%$	1.434 ± 0.004

329 **The trained model is legible, recompilable source.** After training, w^* is written back into a fresh
 330 `.su` as a numeric *literal* with the `w` parameter removed — the trained model expressed as plain Sutra
 331 source:

```
332 function fuzzy rule(vector x, vector own, vector o0, vector o1) {
333     return ((1.431431) * similarity(x, own))
334         && !((1.431431) * similarity(x, o0))
335         && !((1.431431) * similarity(x, o1));
336 }
```

337 This baked `.su` is recompiled through the same PyTorch codegen and, fed the same trained
 338 prototypes, reproduces logits identical to the parametric- w model (max per-logit difference \approx
 339 2×10^{-7} , attributable to float reassociation) — hence identical 100% accuracy. Round-trip ver-
 340 ified for every seed. The trained artifact is therefore not an opaque checkpoint blob: it is a
 341 Sutra program whose learned parameter is a literal in the source, and recompiling that source
 342 reproduces the trained behaviour. Gradient descent here tunes both the embeddings *and* a
 343 weight that survives as readable code — a trained model that is also legible logic. (Two-seed,
 344 single-scale; a verification of the source–training round-trip, not a benchmark.) Reproduction:
 345 `experiments/differentiable_training_weighted.py` (trains `w` + prototypes through the
 346 emitted graph, bakes w^* into `.su`, recompiles, and asserts the round-trip).

347 3.8 Type system and surface syntax — method

348 Sutra's surface syntax is typed: every value carries a primitive class from a fixed set (`int`, `float`,
 349 `complex`, `char`, `bool`, `fuzzy`, `trit`, `vector`, `matrix`, `permutation`, `map`, `string`, `number`,
 350 `void`), and the type drives the synthetic-axis allocation in the extended layout (Appendix B). Type
 351 information is pre-compile-time annotation in the TypeScript sense: it is read by the inliner and the
 352 layout pass before the tensor graph is built, but it is opinionated rather than authoritarian. A diver-
 353 gent assignment warns and still emits a graph, because the runtime guarantee is mathematical not
 354 structural; a type mismatch produces a semantically meaningless but mathematically valid output
 355 rather than a runtime exception. The surface itself presents `if` / `while` / `for` / assignment forms
 356 that read imperatively for ergonomic familiarity; the inliner and egglog simplifier (§4) beta-reduce
 357 these into the functional tensor-op core, so the imperative-looking source is a veneer over the same
 358 compiled graph a hand-written functional spec would produce.

359 A concrete source example grounds the surface. The following is the encode/decode core of
 360 `examples/role_filler_record.su` verbatim — a role-filler record encoded as one vector and a
 361 field decoded back, with no control flow in the program text:

```

362 function vector make_record(vector name, vector color, vector shape) {
363     return bundle(
364         bind(r_name, name),
365         bind(r_color, color),
366         bind(r_shape, shape)
367     );
368 }
369
370 function string decode_field(vector record, vector role) {
371     vector recovered = unbind(role, record);
372     vector winner = argmax_cosine(
373         recovered,
374         [f_alice, f_bob, f_red, f_blue, f_circle, f_square]
375     );
376     return FILLER_NAME[winner];
377 }

```

378 `make_record` beta-reduces to one fused tensor expression (each `bind` a constant-matrix `matmul`,
379 `bundle` a normalized sum); `decode_field` lowers to a single `unbind` `matmul` plus an `argmax-`
380 `cosine` against the codebook. No `bind/bundle/unbind` symbol or host control flow survives in the
381 compiled graph; Appendix F traces an analogous reduction end-to-end.

382

383 4 The Sutra Compiler

384 The compiler is a five-stage pipeline:

- 385 1. **Lex + parse**: `.su` source \rightarrow AST.
- 386 2. **Inline + simplify**: `stdlib` operator definitions inlined; an egglog-based simplifier folds
387 equivalent expressions and runs common-subexpression elimination over the algebra.
- 388 3. **Codegen**: AST \rightarrow Python source emitting PyTorch tensor ops. The emitted module in-
389 cludes the runtime class (`_TorchVSA`) as inline source so the artifact is self-contained.
- 390 4. **Compile-time substrate population**: `embed_batch` fetches embeddings for ev-
391 ery string literal; `populate_sutradb` pushes the codebook into SutraDB;
392 `prewarm_rotation_cache` precomputes role rotations.
- 393 5. **Execute**: emitted module loaded; chosen device (CUDA or CPU) initialized at module
394 `import`; `main()` called; result returned.

395 The runtime class is emitted inline rather than imported because the emitted module *is* the substrate-
396 pure tensor-op graph; every compile-time decision (extended-state-vector dimensions, codebook
397 contents, role rotations, SutraDB path, optional `torch.compile`) is baked into the emitted source.
398 Stages 1–4 run at compile time and stage 5 is the runtime forward pass; Figure~3 diagrams the
399 pipeline as a vertical flow with the residual at each stage.

400 4.1 Substrate-purity invariants

401 Three invariants the compiler enforces: (1) every primitive runs on the substrate (numpy is allowed
402 only at compile time for codebook construction and rotation pre-warm, never on the runtime hot
403 path); (2) no scalar extraction inside an operation: operations may not unpack a Python float from
404 a substrate vector, do scalar arithmetic, and pack the result back; (3) no Python control flow inside
405 an operation: loop halt uses substrate primitives (`heaviside`, `saturate_unit`) instead of Python
406 ternaries.

407 4.2 Compile-time resolution of role rotations

408 The central compile-time mechanism that lets the compiler emit a substrate-pure tensor-op
409 graph is **precomputed rotation matrices**: every role rotation is constructed at compile time
410 (`prewarm_rotation_cache`) and stored as a constant tensor. At runtime, `bind(role, filler)`

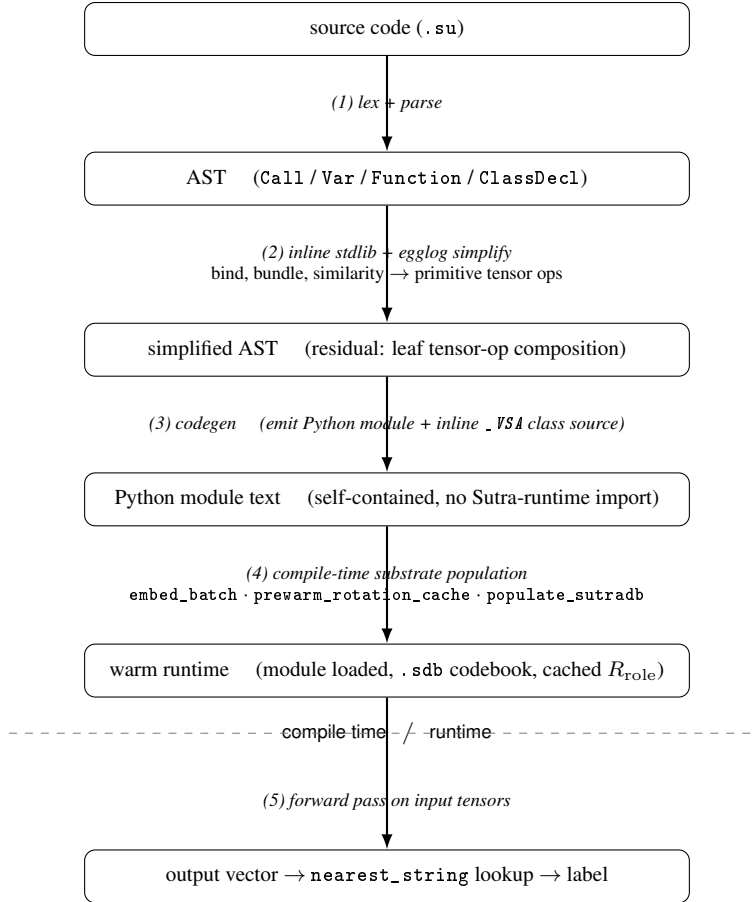


Figure 3: Five-stage compilation pipeline (§4). Boxes are intermediate artifacts; italic labels are the compiler passes that connect them.

411 is a single matmul against a precomputed matrix, the compile-time resolution eliminates the QR construction from the runtime graph entirely. Role rotations are runtime constants, like neural-network
 412 weights at inference; opt-in `torch.compile` (`SUTRA_TORCH_COMPILE=1`) further folds the per-tick
 413 loop body into a single fused kernel. Appendix F traces the lowering of `encode2(r_a, f_a, r_b,`
 414 `f_b) := bundle(bind(r_a, f_a), bind(r_b, f_b))` through every reduction stage; the bot-
 415 tom of the chain contains no `bind/bundle/normalize` symbol and no Python control flow, so
 416 surface lambda calculus and runtime tensor arithmetic are two notations for the same computation.
 417

418

419 5 Demonstration, limitations, and future work

420 The smoke test (`examples/_smoke_test.py`) runs 10 demonstration programs end-to-end across
 421 27 `.su` files (Appendix I); loop coverage lives in `examples/do_while_adder.su` and the 23-case
 422 `test_loop_function_decl.py` suite, and the §3.6 differentiable-training experiment uses the
 423 same primitive set. The embedded codebook covers the compile-time `embed` → runtime `decode`
 424 path; extended features (hashmap routing, persistent codebook via `SUTRA_DB_PATH`) are deferred
 425 pending a concrete requirement.

426 5.1 Limitations

427 The demonstrated regime is bounded, and we state the bounds explicitly:

- 428 • **Composition depth.** Decode is exact for single-cycle records but degrades with chained
429 bind/unbind through bundled distractors: 100% through chain length $L = 2$, falling to
430 chance by $L = 8$ on every substrate (§3.2.1). The demonstrated programs are single-cycle;
431 deep nested records are out of the validated regime.
- 432 • **Substrate dependence.** Every number is measured on four specific frozen substrates.
433 Large-bundle-width decode capacity varies substantially by substrate (§3.2 table) and is
434 not guaranteed for an arbitrary embedding model; behaviour under model drift is future
435 work.
- 436 • **Codebook scale.** The compile-time codebook is $O(\text{vocabulary})$; very large codebooks
437 would require approximate-nearest-neighbour trade-offs not explored here.
- 438 • **No external-system benchmark.** We measure rotation versus Hadamard binding and
439 gradient flow through the compiled graph; we do not benchmark against other neuro-
440 symbolic systems (Scallop, DeepProbLog) on a shared task, and we do not benchmark
441 compiler/runtime wall-clock beyond the indicative timings in Appendix H.
- 442 • **Binding is fixed, not learned.** Role bindings are content-hash-seeded Haar rotations;
443 learned or semantic binding operators are not part of this work.
- 444 • **Training is in-sample.** The §3.6 experiment verifies gradient flow to every learnable pa-
445 rameter, not generalization; no held-out split is reported.

446

447 6 Conclusion

448 Sutra compiles a typed pure-functional source language to a substrate-pure PyTorch tensor-op graph:
449 one vector layout per value, one tensor op per primitive, one dataflow graph per program, no type
450 dispatch at the leaves. With the language in hand, asking which embedding operations compose at
451 what capacity on which substrates becomes a program to write.

452

453 7 Reproducibility

454 Sutra — the language, compiler, runtime, and every .su program cited in this paper — is openly
455 available at <https://github.com/EmmaLeonhart/Sutra>. A self-contained replication pack-
456 age is published at <https://sutra.emmaleonhart.com/sutra-replication-package.zip>:
457 it ships SKILL.md, an agent-runnable recipe an autonomous agent can follow to install the toolchain
458 and re-derive every result reported in this paper end-to-end, with no human in the loop. The
459 project site (this paper in HTML and the conceptual documentation) is at [https://sutra.](https://sutra.emmaleonhart.com)
460 [emmaleonhart.com](https://sutra.emmaleonhart.com); the paper PDF is at <https://sutra.emmaleonhart.com/paper.pdf>.

461

462 8 AI-use statement

463 This work was developed in substantial collaboration with large language models, including for
464 ideation, exploration of the vector-symbolic literature, and drafting. The author independently de-
465 signed and implemented the compiler and runtime, verified that every operation executes on the
466 substrate, ran and checked all experiments, and is responsible for the correctness of every claim,
467 number, and citation in this paper. No results or references were accepted from a model without
468 verification. No experimental result, table, or numerical value reported in this paper was generated
469 by a language model; every number comes from the author-run experiments and reproduction scripts
470 cited in the text and appendices.

471

- 473 • Darwiche, A., & Marquis, P. (2002). A knowledge compilation map. *JAIR* 17:229–264.
- 474 • Gayler, R. W. (2003). Vector symbolic architectures answer Jackendoff's challenges for
- 475 cognitive neuroscience. *Joint International Conference on Cognitive Science*.
- 476 • Kanerva, P. (2009). Hyperdimensional computing: An introduction to computing in dis-
- 477 tributed representation with high-dimensional random vectors. *Cognitive Computation*
- 478 1(2):139–159.
- 479 • Kleene, S. C. (1952). *Introduction to Metamathematics*. North- Holland. The strong three-
- 480 valued logic system used as the ground for Sutra's polynomial fuzzy connectives (§1.1-1).
- 481 • Badreddine, S., Garcez, A. d., Serafini, L., & Spranger, M. (2022). Logic Tensor Networks.
- 482 *Artificial Intelligence* 303:103649.
- 483 • Hájek, P. (1998). *Metamathematics of Fuzzy Logic*. Trends in Logic vol. 4. Kluwer
- 484 Academic. The standard reference for t-norm-based fuzzy logics (Gödel, Łukasiewicz,
- 485 product) cited in §1.1-1 to place Sutra's polynomial connectives.
- 486 • Heddes, M., Nunes, I., Vergés, P., Kleyko, D., Abraham, D., Givargis, T., Nicolau, A.,
- 487 & Veidenbaum, A. (2023). Torchhd: An open source python library to support research
- 488 on hyperdimensional computing and vector symbolic architectures. *Journal of Machine*
- 489 *Learning Research* 24(255):1–10.
- 490 • Li, Z., Huang, J., & Naik, M. (2023). Scallop: A Language for Neurosymbolic Pro-
- 491 gramming. *Proceedings of the ACM on Programming Languages* 7(PLDI):1463–1487.
- 492 arXiv:2304.04812.
- 493 • Manhaeve, R., Dumancic, S., Kimmig, A., Demeester, T., & De Raedt, L. (2018). Deep-
- 494 ProbLog: Neural Probabilistic Logic Programming. *NeurIPS*.
- 495 • Serafini, L. & Garcez, A. d. (2016). Logic Tensor Networks: Deep Learning and Logical
- 496 Reasoning from Data and Knowledge. *NeSy Workshop*.
- 497 • van Krieken, E., Acar, E., & van Harmelen, F. (2022). Analyzing Differentiable Fuzzy
- 498 Logic Operators. *Artificial Intelligence* 302:103602. The differentiable-fuzzy-logic sur-
- 499 vey cited in §1.1-1; analyzes t-norm-derived AND/OR/IMPLIES operators in the neural-
- 500 symbolic context and is the closest prior literature to Sutra's polynomial approach.
- 501 • Vergés, P., Heddes, M., Nunes, I., Givargis, T., & Nicolau, A. (2023). HDCC: A Hy-
- 502 perdimensional Computing compiler for classification on embedded systems and high-
- 503 performance computing. arXiv:2304.12398.
- 504 • Yang, Z., Ishay, A., & Lee, J. (2020). NeurASP: Embracing Neural Networks into Answer
- 505 Set Programming. *IJCAI*.
- 506 • Plate, T. A. (1995). Holographic reduced representations. *IEEE Transactions on Neural*
- 507 *Networks* 6(3):623–641.
- 508 • Siegelmann, H. T. & Sontag, E. D. (1992). On the computational power of neural nets.
- 509 *COLT '92*. Establishes that recurrent neural networks with rational weights are Turing-
- 510 complete; Sutra's tail-recursive loops over a fixed-width state vector take that recurrent
- 511 form.
- 512 • Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic
- 513 structures in connectionist systems. *Artificial Intelligence* 46(1–2):159–216.
- 514 • Ethayarajh, K. (2019). How Contextual are Contextualized Word Representations? Compar-
- 515 ing the Geometry of BERT, ELMo, and GPT-2 Embeddings. *EMNLP-IJCNLP*.
- 516 arXiv:1909.00512. The anisotropy / narrow-cone result: contextual embeddings occupy
- 517 a narrow cone, inflating cosine similarity between unrelated items — the precise reason
- 518 frozen LLM embeddings are not the i.i.d. distribution textbook VSA assumes (§1, §2).
- 519 • Gao, J., He, D., Tan, X., Qin, T., Wang, L., & Liu, T.-Y. (2019). Representation Degenera-
- 520 tion Problem in Training Natural Language Generation Models. *ICLR*. arXiv:1907.12009.
- 521 • Mu, J., Bhat, S., & Viswanath, P. (2018). All-but-the-Top: Simple and Effective Postpro-
- 522 cessing for Word Representations. *ICLR*. arXiv:1702.01417.
- 523 • Kingma, D. P. & Ba, J. (2015). Adam: A Method for Stochastic Optimization. *3rd Interna-*
- 524 *tional Conference on Learning Representations (ICLR)*. arXiv:1412.6980. The optimizer
- 525 (run with its default $\beta_1, \beta_2, \varepsilon$) used for the §3.6 differentiable-training experiment.
- 526 • Lin, Z., Akin, H., Rao, R., Hie, B., Zhu, Z., Lu, W., Smetanin, N., Verkuil, R., Kabeli,
- 527 O., Shmueli, Y., dos Santos Costa, A., Fazel-Zarandi, M., Sercu, T., Candido, S., & Rives,
- 528 A. (2023). Evolutionary-scale prediction of atomic-level protein structure with a language

529

model. *Science* 379(6637):1123–1130. The ESM-2 protein language model used as the non-text substrate in §3.2.

530

531

- Zadeh, L. A. (1965). Fuzzy sets. *Information and Control* 8(3):338–353.

532

- Jang, J.-S. R. (1993). ANFIS: Adaptive-Network-Based Fuzzy Inference System. *IEEE Transactions on Systems, Man, and Cybernetics* 23(3):665–685.

533

534

- Buckley, J. J. & Hayashi, Y. (1994). Fuzzy neural networks: A survey. *Fuzzy Sets and Systems* 66(1):1–13.

535

536

537 **Appendix**

538 **Appendix A. Notation: extended layout and primitive operations**

539 We work in a fixed-dimensional real vector space \mathbb{R}^d where d is the substrate's embedding dimension
 540 (768 for nomic-embed-text, 384 for all-minilm, 1024 for mxbai-embed-large, 320 for ESM-2). Ev-
 541 ery Sutra value carries the extended layout [semantic | synthetic], a d -dimensional semantic block
 542 holding the substrate embedding, concatenated with a small fixed-width synthetic block reserving
 543 canonical axes for primitive types (real, imag, truth, char, loop-done) and slot machinery (§3.3).
 544 Where notation does not distinguish, "vector" means "the full extended-layout tensor."

545 The seven primitive operations are:

$$\begin{aligned} \text{bind}(r, f) &= R_r f, & R_r &= \text{QR}(\text{seed} = \text{hash}(r)).Q \\ \text{unbind}(r, v) &= R_r^\top v \\ \text{bundle}(x, y) &= \frac{x + y}{\|x + y\| + \varepsilon} \\ \text{similarity}(x, y) &= \frac{x \cdot y}{\|x\| \|y\| + \varepsilon} \\ \text{normalize}(v) &= \frac{v}{\|v\| + \varepsilon} \end{aligned}$$

546 plus the Lagrange Kleene gates (scalar \rightarrow scalar, exact on the $\{-1, 0, +1\}^2$ grid, §1.1-1) and the
 547 soft-halt cell (state, halt \rightarrow state', halt', §3.4).

548 The Lagrange gates in closed form:

$$\begin{aligned} \text{AND}(a, b) &= \frac{1}{2}(a + b + ab - a^2 - b^2 + a^2b^2) \\ \text{NAND}(a, b) &= \frac{1}{2}(-a - b - ab + a^2 + b^2 - a^2b^2) \\ \text{OR}(a, b) &= \frac{1}{2}(a + b - ab + a^2 + b^2 - a^2b^2) \\ \text{NOR}(a, b) &= \frac{1}{2}(-a - b + ab - a^2 - b^2 + a^2b^2) \\ \text{NOT}(a) &= -a \\ \text{XOR}(a, b) &= -ab \\ \text{XNOR}(a, b) &= ab \end{aligned}$$

549 The soft-halt cell update is, in compact form,

$$\begin{aligned} s_{t+1} &= R s_t && \text{(rotation step)} \\ h_t &= \text{Heaviside}(\text{cond}(s_t)) && \text{(per-tick halt signal)} \\ H_t &= \text{sat}_{[0,1]} \left(\sum_{k \leq t} h_k \right) && \text{(cumulative monotone halt)} \\ \hat{s}_{t+1} &= H_t s_t + (1 - H_t) s_{t+1} && \text{(soft-mux freeze)} \end{aligned}$$

550 Every right-hand side is a tensor expression with no Python control flow. The compile-time primi-
 551 tives `RotationFor` and `embed` produce constants R_r and basis vectors at compile time and are not
 552 part of the runtime tensor graph.

553 **Appendix B. Extended-state-vector layout: per-axis assignments**

554 §3.3 describes the [semantic | synthetic] layout in prose. The diagram and per-axis purpose
 555 table below give the concrete allocation referenced in `codegen_pytorch.py`:

```

556      +-----+-----+-----+-----+-----+-----+
557  value | semantic block      | R | I | T | C | L | slots... |
558      +-----+-----+-----+-----+-----+-----+
559      |<-- semantic_dim ----->|<-- synthetic_dim ----->|>
560      0  1  2  3  4  5..
561      REAL IMAG TRUTH CHAR LOOP_DONE
562      _FLAG

```

Index	Purpose
synthetic[0]	AXIS_REAL (real component for int/float/complex)
synthetic[1]	AXIS_IMAG (imaginary component for complex)
synthetic[2]	AXIS_TRUTH (fuzzy truth scalar; bool/comparisons)
synthetic[3]	AXIS_CHAR_FLAG (marks char primitives)
synthetic[4]	AXIS_LOOP_DONE (substrate-side completion flag)
synthetic[5..]	SLOT_BASE: disjoint 2D Givens slots for variable storage

563 At semantic_dim = 768 (nomic-embed-text), synthetic_dim = 100 accommodates the five
564 canonical axes plus 47 disjoint Givens slots.

565 **Appendix C. Capacity: full per-substrate sweeps**

566 Cross-substrate decode accuracy at full bundle widths $k \in \{2, 4, 8, 16, 24, 32, 48\}$. The four
567 substrates use 84-entry vocabularies (LLM substrates: 84-word noun set spanning animals, foods,
568 objects, places, abstract nouns; ESM-2: 84-sequence amino-acid set covering canonical signal pep-
569 tides, cell-penetrating peptides, antimicrobial peptides, classic affinity-tag motifs, and deterministic
570 random k-mers). All embeddings are unit-normalized; nomic-embed-text and ESM-2 are addition-
571 ally mean-centered.

572 **nomic-embed-text (768-d, mean-centered):**

k	rotation accuracy	rotation signal cos	Hadamard accuracy	Hadamard signal cos
2	100.0%	+0.703	95.0%	+0.488
4	100.0%	+0.497	95.0%	+0.400
8	100.0%	+0.354	87.5%	+0.307
16	100.0%	+0.251	84.4%	+0.230
24	100.0%	+0.203	60.8%	+0.189
32	99.1%	+0.176	63.1%	+0.167
48	93.3%	+0.144	48.3%	+0.136

573 **all-minilm (384-d):**

k	rotation accuracy	rotation signal cos	Hadamard accuracy	Hadamard signal cos
2	100.0%	+0.711	45.0%	+0.386
4	100.0%	+0.506	10.0%	+0.335
8	100.0%	+0.356	7.5%	+0.315
16	92.5%	+0.252	3.1%	+0.299
24	76.2%	+0.203	2.9%	+0.300
32	66.9%	+0.179	2.5%	+0.297
48	42.3%	+0.144	1.7%	+0.294

574 **mxbai-embed-large (1024-d):**

k	rotation accuracy	rotation signal cos	Hadamard accuracy	Hadamard signal cos
2	100.0%	+0.708	15.0%	+0.311
4	100.0%	+0.500	2.5%	+0.304
8	100.0%	+0.353	2.5%	+0.295
16	98.8%	+0.251	1.2%	+0.294
24	95.8%	+0.203	0.8%	+0.293
32	85.3%	+0.176	0.9%	+0.292
48	72.1%	+0.146	1.0%	+0.291

575 **ESM-2 small protein language model (320-d, mean-centered):**

k	rotation accuracy	rotation signal cos	Hadamard accuracy	Hadamard signal cos
2	100.0%	+0.713	75.0%	+0.470
4	100.0%	+0.501	50.0%	+0.323
8	100.0%	+0.349	28.7%	+0.257
16	90.6%	+0.252	16.2%	+0.185
24	77.1%	+0.205	11.2%	+0.171
32	61.9%	+0.174	6.2%	+0.141
48	44.2%	+0.143	4.2%	+0.117

576 The signal cosine for Hadamard is comparable to rotation's, but the noise floor is much higher be-
577 cause the elementwise product of correlated real-valued embeddings produces a result that overlaps
578 with many distractors in the codebook rather than near-orthogonally with one.

579 **Appendix D. Crosstalk depth: full per-substrate L-sweep**

580 The §3.2.1 protocol: chain length $L \in \{1, 2, 4, 8, 16, 32\}$, 20 trials, bundle width 4 (3 distractors
581 per cycle). Forward-bind through L role rotations bundling 3 distractor (role, filler) pairs at each
582 step; unbind in reverse and decode. Two flavors: *raw* (no cleanup) and *snap* (argmax-cosine cleanup
583 against the codebook after each unbind step).

substrate	L=1 raw	L=2 raw	L=4 raw	L=1 snap	L=2 snap	L=4 snap
nomic-embed-text	100%	100%	20%	100%	10%	0%
all-minilm	100%	100%	5%	100%	0%	0%
mxbai-embed-large	100%	100%	5%	100%	0%	0%

584 By chain length 8 raw accuracy is at chance (1/84) on all three substrates. Snap is *worse* than
585 raw past chain length 1: a hard codebook commitment converts soft noise into a high-confidence
586 wrong answer that the next unbind cannot recover from. The runtime does not implicitly snap
587 between operations; cleanup is an explicit step the program schedules where it knows the codebook
588 is the right reference. Reproduction script: `experiments/crosstalk_chain.py`; raw JSON in
589 `experiments/crosstalk_chain_results.json`.

590 **Appendix E. Codebook implementation details**

591 The §3.5 codebook is implemented as an embedded vector database (internally SutraDB) shipped
592 as part of the compiler, analogous to SQLite being embedded in an application rather than run as
593 a separate service. The data model is RDF triples with f32-vector literals as the object position,
594 indexed by a built-in HNSW index for nearest-neighbor decode. The on-disk format is a `.sdb` file
595 that travels alongside the compiled Python module; no external service, no separate install, no net-
596 work dependency. Every embedded string in a Sutra program is inserted with the embedding as the
597 object of a triple typed `<http://sutra.dev/f32vec>`. Strings declared but unused in expressions

598 are still inserted, so they remain decodable. The compiled module's Python data section never carries the embeddings, they live in the .sdb file, an artifact of compilation, not a service the runtime contacts.
600

601 nearest_string runs over an HNSW (Hierarchical Navigable Small World) approximate-nearest-neighbor graph maintained by the triplestore. HNSW (Malkov & Yashunin, TPAMI 2020) has **O(log N) expected and worst-case query time** under standard graph-construction parameters; it has displaced linear scan as the default ANN index in Faiss, Milvus, Weaviate, Qdrant, and most production vector databases. A 100-string codebook and a 100,000-string codebook have comparable decode latency at runtime, modulo HNSW's tunable M (graph degree) and ef_search (beam width); the cost difference is roughly one extra graph hop per $10\times$ growth in N.
607

608 Appendix F. Worked lowering of a two-field bundled record

609 The body §4.2 sketches the lowering of $\text{encode2}(r_a, f_a, r_b, f_b) := \text{bundle}(\text{bind}(r_a, f_a), \text{bind}(r_b, f_b))$. Here we trace each stage with the explicit residual.
610

611 **Stage 1: AST after parse.** A tree of Call nodes over named identifiers: $\text{Call}(\text{"bundle"}, \text{Call}(\text{"bind"}, r_a, f_a), \text{Call}(\text{"bind"}, r_b, f_b))$.
612

613 **Stage 2: beta reduction by stdlib inlining.** bind, bundle, and normalize are stdlib functions: $\text{bind}(r, f) \equiv \text{RotationFor}(r) f$, $\text{bundle}(x, y) \equiv \text{normalize}(x + y)$, $\text{normalize}(v) \equiv v / (\|v\| + \varepsilon)$.
614 After substitution the body becomes
615

$$\text{normalize}(\text{RotationFor}(r_a) f_a + \text{RotationFor}(r_b) f_b).$$

616 No bind or bundle symbol remains; the residual is straight-line algebra over four tensor primitives.

617 **Stage 3: compile-time constant resolution.** $\text{RotationFor}(r)$ is a compile-time function returning $R = \text{QR}(\text{seed} = \text{hash}(r)).Q$. The compiler evaluates it for each role at compile time, freezes the results as constant tensors R_a and R_b , and stores them in the rotation cache. The body becomes
618 $\text{normalize}(R_a f_a + R_b f_b)$, R_a and R_b are now load-bearing constants in the same sense as the weight matrices of a feed-forward network.
619
620
621

622 **Stage 4: peephole fusion.** The simplifier recognizes $\text{normalize}(\sum_i R_i f_i)$ as the bundle-of-binds pattern and rewrites it to $\text{_VSA.bundle_of_binds}([(R_a, f_a), (R_b, f_b)])$, one kernel launch instead of two matmuls + add + norm.
623
624

625 **Stage 5: leaf tensor ops at runtime.** bundle_of_binds stacks rotations into a (k, d, d) tensor, stacks fillers into (k, d) , runs one batched einsum + sum + L2-normalize:
626

$$v = \sum_k R_k f_k = \text{einsum}(\text{"kij, kj- > i"}, \text{stack}([R_a, R_b]), \text{stack}([f_a, f_b]))$$

$$\text{encode2} = v / (\|v\| + \varepsilon)$$

627 The compiled forward pass for encode2 is exactly those three torch calls (einsum, linalg.norm, divide) over precomputed R_a, R_b and runtime-supplied f_a, f_b .
628

629 Appendix G. §3.6 differentiable-training vocabulary

630 Twenty categories of fifty words each (992 unique after deduplication), embedded via nomic-embed-text:
631

- 632 • **animal:** dog, cat, bird, fish, horse, lion, tiger, elephant, rabbit, monkey, bear, wolf, fox, deer, mouse, snake, frog, turtle, dolphin, whale, shark, eagle, owl, sparrow, crow, robin, parrot, swan, duck, goose, chicken, cow, pig, sheep, goat, donkey, camel, giraffe, kangaroo, koala, panda, leopard, cheetah, hippopotamus, rhinoceros, antelope, buffalo, hedgehog, squirrel, raccoon
- 633
- 634
- 635
- 636
- 637 • **vehicle:** car, truck, airplane, boat, bicycle, motorcycle, bus, train, ship, helicopter, tractor, scooter, van, taxi, jeep, sailboat, kayak, canoe, raft, submarine, glider, jet, rocket, space-ship, sled, skateboard, wagon, carriage, chariot, ambulance, firetruck, limousine, minivan,
- 638
- 639

640 hatchback, sedan, coupe, convertible, pickup, trailer, ferry, yacht, dinghy, blimp, balloon,
641 hovercraft, tram, moped, tricycle, rollerblade, unicycle

- 642 • **food:** apple, bread, cheese, rice, pasta, banana, salad, soup, meat, pizza, sandwich, burger,
643 taco, sushi, cake, cookie, pie, donut, muffin, pancake, waffle, bagel, croissant, omelet,
644 salmon, tuna, beef, pork, lamb, bacon, ham, sausage, steak, lobster, shrimp, crab, oyster,
645 clam, broccoli, carrot, lettuce, tomato, potato, cucumber, onion, garlic, pepper, eggplant,
646 spinach, mushroom
- 647 • **color:** red, blue, green, yellow, orange, purple, black, white, brown, pink, gray, cyan,
648 magenta, violet, indigo, turquoise, teal, lavender, maroon, crimson, scarlet, ruby, gold,
649 silver, bronze, copper, beige, tan, ivory, charcoal, navy, sapphire, emerald, jade, olive,
650 lime, mint, coral, peach, plum, mauve, fuchsia, amber, ochre, sienna, mahogany, chocolate,
651 caramel, mustard, azure
- 652 • **clothing:** shirt, pants, dress, hat, shoes, jacket, socks, gloves, scarf, belt, sweater, hoodie,
653 jeans, shorts, skirt, blouse, coat, cap, beanie, mittens, tights, leggings, vest, blazer, suit,
654 tuxedo, gown, robe, kimono, kilt, poncho, cloak, cape, sneakers, boots, sandals, slip-
655 pers, heels, loafers, tie, bowtie, cufflinks, watch, ring, necklace, earrings, bracelet, anklet,
656 brooch, headband
- 657 • **weather:** rain, snow, wind, cloud, storm, fog, frost, hail, thunder, lightning, drizzle, down-
658 pour, blizzard, hurricane, tornado, cyclone, typhoon, sleet, mist, haze, smog, sunshine,
659 sunlight, sunset, sunrise, dawn, dusk, twilight, breeze, gust, gale, humidity, drought, flood,
660 monsoon, snowfall, snowstorm, rainstorm, sandstorm, heatwave, chill, dew, hailstorm,
661 thaw, overcast, sunny, cloudy, rainy, snowy, windy
- 662 • **emotion:** joy, sadness, anger, fear, love, hope, surprise, disgust, pride, envy, happiness,
663 grief, rage, anxiety, affection, despair, delight, shame, guilt, confidence, contentment, jeal-
664 ously, regret, sorrow, frustration, satisfaction, awe, wonder, gratitude, compassion, sym-
665 pathy, empathy, irritation, boredom, excitement, enthusiasm, calm, serenity, melancholy,
666 nostalgia, longing, embarrassment, humiliation, indifference, ecstasy, bliss, dread, terror,
667 amusement, loneliness
- 668 • **tool:** hammer, saw, drill, wrench, screwdriver, knife, scissors, pliers, axe, shovel, rake, hoe,
669 spade, pickaxe, crowbar, mallet, chisel, sander, level, ruler, vise, clamp, ratchet, socket,
670 awl, scraper, trowel, broom, mop, sponge, bucket, ladder, jackhammer, sledgehammer,
671 paintbrush, roller, stapler, tongs, tweezers, calipers, magnifier, flashlight, multimeter, wire-
672 cutter, hacksaw, router, torch, soldering_iron, drillbit, screwbit
- 673 • **instrument:** guitar, piano, drum, violin, flute, trumpet, saxophone, harp, cello, clarinet,
674 banjo, mandolin, ukulele, harmonica, accordion, organ, keyboard, synthesizer, xylophone,
675 tambourine, maracas, bongos, marimba, vibraphone, glockenspiel, bagpipes, oboe, bas-
676 soon, trombone, tuba, lute, sitar, koto, zither, dulcimer, cymbal, gong, triangle, cowbell,
677 snare, kettledrum, recorder, piccolo, fife, didgeridoo, theremin, viola, double_bass, fiddle,
678 ocarina
- 679 • **profession:** doctor, teacher, lawyer, engineer, nurse, chef, artist, scientist, farmer, plumber,
680 electrician, carpenter, mechanic, pilot, sailor, soldier, judge, journalist, writer, poet, painter,
681 sculptor, musician, actor, dancer, singer, photographer, architect, dentist, surgeon, pharma-
682 cist, veterinarian, librarian, accountant, banker, broker, programmer, designer, manager,
683 secretary, butcher, baker, gardener, tailor, jeweler, barber, chemist, biologist, physicist,
684 mathematician
- 685 • **body_part:** head, hand, foot, eye, ear, nose, mouth, leg, arm, finger, toe, knee, elbow,
686 shoulder, hip, neck, back, chest, stomach, heart, brain, lung, liver, kidney, bone, muscle,
687 skin, hair, throat, jaw, chin, cheek, forehead, eyebrow, eyelash, lip, tongue, palm, wrist,
688 ankle, thumb, heel, spine, rib, scalp, nostril, gum, knuckle, tendon, vein
- 689 • **plant:** tree, flower, grass, bush, vine, fern, moss, herb, weed, leaf, stem, branch, bark,
690 blossom, petal, oak, maple, willow, birch, cedar, bamboo, cactus, rose, tulip, daisy, lily,
691 sunflower, orchid, ivy, basil, rosemary, thyme, sage, lavender, dandelion, clover, lotus,
692 magnolia, sycamore, redwood, baobab, eucalyptus, juniper, hemlock, fir, spruce, ash, elm,
693 poplar, chestnut
- 694 • **furniture:** chair, table, sofa, bed, desk, shelf, drawer, cabinet, wardrobe, dresser, night-
695 stand, ottoman, bench, stool, recliner, futon, couch, armchair, bookcase, sideboard, buffet,
696 cupboard, hutch, vanity, headboard, footboard, mattress, pillow, cushion, blanket, quilt,
697 comforter, lamp, mirror, rug, carpet, curtain, blind, shutter, hammock, cradle, crib, bassinet,
698 highchair, rocker, loveseat, settee, divan, chaise, headrest

- 699 • **building**: house, apartment, mansion, cottage, cabin, hut, igloo, tent, palace, castle,
700 fortress, tower, skyscraper, office, factory, warehouse, store, mall, restaurant, hotel, mo-
701 tel, hospital, school, university, library, museum, theater, stadium, arena, church, temple,
702 mosque, synagogue, cathedral, chapel, monastery, abbey, barn, shed, garage, basement,
703 attic, cellar, lobby, lounge, hallway, corridor, atrium, foyer, balcony
- 704 • **country**: France, Germany, Italy, Spain, Portugal, England, Scotland, Ireland, Norway,
705 Sweden, Finland, Denmark, Iceland, Russia, Poland, Greece, Turkey, Egypt, Morocco, Al-
706 geria, Kenya, Nigeria, Ethiopia, Ghana, Senegal, Mali, Sudan, Uganda, Tanzania, Mada-
707 gascar, China, Japan, Korea, Vietnam, Thailand, Malaysia, Indonesia, India, Pakistan,
708 Bangladesh, Iran, Iraq, Israel, Lebanon, Australia, Canada, Mexico, Brazil, Argentina,
709 Chile
- 710 • **sport**: football, basketball, baseball, soccer, tennis, golf, hockey, rugby, cricket, volley-
711 ball, swimming, running, cycling, skiing, snowboarding, surfing, sailing, rowing, kayaking,
712 climbing, hiking, boxing, wrestling, fencing, archery, shooting, fishing, hunting, polo, bad-
713 minton, ping_pong, squash, racquetball, lacrosse, handball, dodgeball, kickball, gymnas-
714 tics, diving, weightlifting, judo, karate, taekwondo, sumo, marathon, triathlon, decathlon,
715 biathlon, skating, bowling
- 716 • **drink**: water, juice, milk, tea, coffee, soda, beer, wine, whiskey, vodka, rum, gin, tequila,
717 brandy, cognac, champagne, cocktail, smoothie, milkshake, lemonade, cider, ale, lager,
718 stout, bourbon, scotch, sake, mead, punch, eggnog, kombucha, kefir, espresso, latte, cap-
719 puccino, mocha, americano, macchiato, frappe, hot_chocolate, cordial, shake, slushie,
720 syrup, fizz, brew, tonic, infusion, ginger_ale, root_beer
- 721 • **metal**: gold, silver, copper, iron, steel, aluminum, brass, bronze, tin, lead, zinc, nickel, plat-
722 inum, titanium, chromium, mercury, magnesium, lithium, sodium, potassium, calcium, ura-
723 nium, plutonium, palladium, tungsten, vanadium, cobalt, manganese, beryllium, gallium,
724 indium, antimony, bismuth, cadmium, cerium, neodymium, osmium, rhodium, ruthenium,
725 tantalum, thallium, thorium, yttrium, scandium, hafnium, niobium, molybdenum, rhenium,
726 iridium, rubidium
- 727 • **shape**: circle, square, triangle, rectangle, oval, ellipse, pentagon, hexagon, octagon, dia-
728 mond, rhombus, trapezoid, parallelogram, polygon, sphere, cube, cylinder, cone, pyramid,
729 prism, cuboid, tetrahedron, dodecahedron, icosahedron, octahedron, torus, helix, spiral,
730 crescent, star, heart, arrow, cross, line, curve, arc, ring, loop, knot, dot, vertex, edge, angle,
731 parabola, hyperbola, sine, wave, zigzag, scallop, annulus
- 732 • **fabric**: cotton, wool, silk, linen, polyester, nylon, denim, leather, suede, velvet, satin, lace,
733 tweed, cashmere, mohair, fleece, fur, canvas, burlap, jute, flannel, chiffon, organza, taffeta,
734 brocade, damask, paisley, gingham, plaid, herringbone, corduroy, microfiber, spandex, ly-
735 cra, rayon, viscose, acrylic, polypropylene, jersey, knit, sherpa, gabardine, twill, muslin,
736 gauze, mesh, vinyl, tulle, georgette, voile

737 Appendix H. Reproduction details and hyperparameters

738 Per-experiment configuration. All scripts live under `experiments/` in the source repository; each
739 writes a JSON results file to the same directory on completion. RNG seeds are fixed in the source
740 files cited; re-running reproduces the numbers reported in the body to the precision reported.

- 741 • **Rotation vs Hadamard, LLM** (§3.2) — script `rotation_binding_capacity_llm.py`;
742 10 trials per k ; substrates `nomic-embed-text`, `all-minilm`, `mxbai-embed-large`; seeds fixed
743 per script.
- 744 • **Rotation vs Hadamard, ESM-2** (§3.2) — script
745 `rotation_binding_capacity_bioinformatics.py`; 10 trials per k ; substrate
746 `facebook/esm2_t6_8M_UR50D`; seeds 1729 and 2718.
- 747 • **Crosstalk depth** (§3.2.1) — script `crosstalk_chain.py`; 20 trials per chain length L ;
748 three LLM substrates; seeds fixed per script.
- 749 • **Differentiable training** (§3.6) — script `differentiable_training_compiled.py`
750 `--batched`; 3 seeds \times 30 epochs ($K=5$, 50 words); substrate `nomic-embed-text` (frozen);
751 optimizer Adam, `lr=0.01`; seeds 0–2.
- 752 • **Trained weight \rightarrow legible source** (§3.7) — script
753 `differentiable_training_weighted.py`; 2 seeds \times 30 epochs ($K=3$, 24 words);
754 substrate `nomic-embed-text` (frozen); optimizer Adam, `lr=0.02`; seeds 0–1.

755 The differentiable-training run (`differentiable_training_compiled.py`) generates a `.su`
 756 fuzzy-rule classifier, compiles it with the PyTorch codegen, and backpropagates five randomly-
 757 initialized unit-normalized prototype vectors through the *emitted* rule function (the compiler's
 758 `_VSA.similarity` composed with the Lagrange–Kleene polynomials). Five classes, ten words
 759 each (50 inputs); embeddings are precomputed once and cached to `.diff_train_embeddings.pt`
 760 so reruns skip the embed step. A build-time assertion checks the emitted similarity is not
 761 `float()`-collapsed (Stage A0), so gradients provably flow through the compiled graph rather than a
 762 reimplement. The forward is evaluated batched via `torch.vmap` over that same emitted rule;
 763 a runtime assertion requires the batched and per-sample evaluations to agree within 10^{-4} before
 764 training, so the fast path is the identical compiled computation, not a substitute.

765 The §3.7 run (`differentiable_training_weighted.py`) adds a number `w` parameter to the `.su`
 766 rule, trains `w` together with the prototypes through the same emitted graph, then regenerates the rule
 767 with `w*` substituted as a numeric literal and the parameter removed, recompiles that source through
 768 the PyTorch codegen, and asserts the recompiled program's logits match the parametric model (max
 769 per-logit difference $< 10^{-4}$) — a source-level training round-trip, not a benchmark.

770 Hardware used for the numbers in the body: CPU torch on a single laptop (no CUDA). The §3.6
 771 compiled run takes ≈ 230 s ($K=5$, 50 words, 30 epochs, 3 seeds) via the batched `torch.vmap` path
 772 over the same emitted ops; the equivalent per-sample Python driver produces the bit-identical result
 773 but takes ≈ 6.2 h — an interpreter-overhead artifact, not a cost of the compiled graph; the §3.7
 774 weighted round-trip ($K=3$, 24 words, 30 epochs, 2 seeds, per-sample — small enough that the driver
 775 cost is ≈ 2.5 min) completes including the recompile check; the §3.2 capacity sweeps complete in
 776 ~ 2 min per substrate; the §3.2.1 crosstalk sweep completes in ~ 5 min. Re-running on CUDA should
 777 reproduce the same accuracy numbers since the operations are deterministic given a seed.

778 Appendix I. Demonstration corpus

779 The smoke test (`examples/_smoke_test.py`) compiles and runs ten `.su` programs end-to-end
 780 and asserts each output against a hardcoded expected value. The programs collectively exercise the
 781 language features the body claims, with no Python control flow on the runtime path:

Program	Feature exercised
<code>hello_world.su</code>	embed + retrieve (minimal program)
<code>fuzzy_branching.su</code>	weighted-superposition conditional
<code>role_filler_record.su</code>	bind / bundle / unbind on a 3-field record (§2.1)
<code>classifier.su</code>	cosine-similarity classifier over a small codebook
<code>analogy.su</code>	associative pair memory: capital \rightarrow country recovery via unbind
<code>knowledge_graph.su</code>	(subject, relation, object) triple encode + decode
<code>predicate_lookup.su</code>	bind-keyed dictionary read
<code>fuzzy_dispatch.su</code>	Lagrange-Kleene-gated dispatch among handlers
<code>nearest_phrase.su</code>	top-1 phrase retrieval over a <code>.sdb</code> codebook
<code>sequence.su</code>	foreach reduction over a list

782 Loop coverage lives in `examples/do_while_adder.su` and the 23-case
 783 `tests/test_loop_function_decl.py` suite. The §3.6 differentiable-training experiment
 784 uses the same primitive set the smoke-test programs are built from, no Sutra-runtime extensions,
 785 just compilation of `.su` source to PyTorch tensor ops.