QUILL: IN-WEIGHT COMPUTATION WITH MLPS

Anonymous authors
Paper under double-blind review

ABSTRACT

Language models continue to improve across mathematics, programming, and long-context reasoning tasks. Nevertheless, elementary arithmetic remains unreliable. Existing remedies use additional reasoning tokens or dedicated numerical executors. We instead study *in-weight computation*: encoding numerical functions directly in ordinary MLP weights for forward-pass execution, with intermediate values retained as neural activations. We introduce **QUILL** (**QU**asi-**I**nterpo**L**ant **L**ayers), which constructs single-hidden-layer MLP approximations using quasi-interpolation theory. For univariate analytic functions, QUILL achieves geometric error decay to a finite-precision floor, with $O(\log(1/\varepsilon))$ width and working precision for target error ε. The construction also reveals conditions under which gradient descent is slow to acquire the hidden-layer scales supporting high precision. We extend QUILL to multivariate polynomial programs by composing square modules into multipliers. For bounded computations in the realizable precision regime, the resulting QUILL circuits attain output error ε with $O(\log(2+\kappa/\varepsilon))$ nonlinear neurons per multiplication, where κ bounds the amplification of local approximation errors. Tested polynomial programs, including dot products and determinants, achieve sampled relative L^2 errors below 6.5×10^{-14} in FP64. Additional primitives support non-polynomial functions, fixed-step numerical algorithms, and bounded discrete programs. As a proof of concept, we install a width-80 QUILL squaring module in three frozen language models. With engineered routing and input/output handling, the installed pathway returns correct answers in all $2{,}250$ evaluations across three-, five-, and seven-digit inputs.

1 INTRODUCTION

Language models continue to improve across mathematics, programming, and long-context reasoning tasks. Frontier systems have reached a gold-medal standard at the International Mathematical Olympiad (Luong et al., 2025). However, current models remain unreliable at elementary arithmetic. Bie et al. (2026) show that every frontier model they evaluate, without execution tools, exhibits substantial degradation on integer addition as the number of digits grows. We observe similar failures in open-weight models. Without tools, Qwen2.5-7B achieves 79.2% exact accuracy on three-digit squaring, falling to 4.8% on seven-digit inputs.

Existing remedies improve arithmetic reliability by adding reasoning steps or offloading the computation outside the network's weights. Chain-of-thought prompting expands the computation across generated intermediate tokens, but adds inference cost without guaranteeing correctness (Wei et al., 2023). Program-aided and tool-augmented models instead rely on an external interpreter or calculator for numerical execution. Recent systems introduce exact executors inside the model, but as dedicated non-neural components. **In contrast, we study in-weight computation**: compiling the numerical function directly into ordinary MLP weights and evaluating it during the forward pass. Its intermediate values and outputs remain ordinary activations, allowing constructed operations to feed directly into learned layers or other constructed modules without intermediate generated tokens or a separate numerical executor.

The first bottleneck to in-weight arithmetic is precision: squaring a seven-digit integer can produce a fourteen-digit answer. Recovering that answer exactly requires relative error below roughly 5×10^{-15}. Gradient training does not approach this high-precision regime: in our experiments, single-hidden-layer MLPs trained with GD or Adam remain over 10 orders of magnitude above

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

the floating point precision floor after millions of updates (Section 3.5). Explicit constructions for single-hidden-layer MLPs with standard smooth activations face a tradeoff between network size and working precision. Classical staircase constructions require $O(1/\varepsilon)$ hidden units to achieve ε error, while constructions with geometric convergence in width use working precision that grows superlinearly in $\log(1/\varepsilon)$.

The second obstacle is stable composition. Computations such as dot products, determinants, and iterative solvers combine many primitive operations, making approximate outputs the inputs to subsequent calculations. Errors propagate according to intermediate magnitudes and sensitivities and can amplify across depth. A realizable construction must therefore control end-to-end error while accounting for how network size and depth grow with the computation.

We address these challenges with QUILL (Quasi-Interpolant Layers), a constructive framework for high-precision computation in ordinary MLP weights. QUILL couples hidden-feature spacing and slope to obtain geometric refinement while keeping the total readout magnitude bounded as width grows, enabling high-precision approximation in standard floating-point arithmetic. The construction also reveals why gradient training is slow to acquire the hidden-layer scales that support high precision. Using squaring to implement multiplication, we realize networks of QUILL modules that can represent in-weight computations ranging from polynomial functions to discrete graph algorithms and fixed-step numerical procedures. Applying arithmetic circuit theory, we derive explicit bounds on end-to-end error and network complexity. Finally, as a proof of concept for in-weight execution, we install a constructed squaring module into frozen pretrained language models and execute it during the forward pass.

Our contributions are:

- **QUILL: High-precision MLP primitives (Section 3):** We construct single-hidden-layer MLPs for bounded analytic functions with geometric refinement in width down towards a finite-precision floor. The construction combines $O(\log(1/\varepsilon))$ width and working precision with a width-independent readout ℓ_1 bound (Theorem 3.1, Corollary 3.1.1), ensuring floating point realizability. QUILL reveals why standard training is unable to reach high precision, thus justifying our construction approach. Informally, high-precision MLPs require large pre-activation slopes; we prove that under sufficiently dispersed gradient signal across neurons, gradient descent is unable to learn the correct pre-activation slope magnitudes fast enough.
- **QUILL circuits (Section 4):** We compose QUILL primitives into numerical and discrete programs, including determinants, softmax, fixed-step Newton iteration, and shortest paths. For bounded arithmetic computations, we derive end-to-end error bounds and bounds on the nonlinear units and layers needed to achieve a target accuracy (Theorem 4.1). In FP64, sampled relative L^2 error stays below 6.5×10^{-14} across the polynomial suite; sorting and shortest paths return exact decoded outputs on the evaluated instances.
- **LLM Integration (Section 5):** As a proof of concept, we show an example of installing QUILL circuits into LLMs. Installed in three frozen LMs with hard-coded routing, a width-80 QUILL square module answers all $2{,}250$ evaluated 3-, 5-, and 7-digit squaring problems with 100% accuracy.

2 Related Work

Numerical computation in language models. Existing approaches differ largely in where numerical execution occurs. Chain-of-thought performs computation through generated intermediate reasoning steps (Wei et al., 2023), while Program-of-Thoughts, PAL, and Toolformer delegate it to external interpreters or APIs (Chen et al., 2023; Gao et al., 2023; Schick et al., 2023). Recent work instead moves execution inside the model through specialized mechanisms. Dugan et al. (2024) couple a frozen LLM to a symbolic executor, while Dietz & Klakow (2025) integrate a gated calculator directly into the model. We instead encode numerical functions directly into ordinary MLP weights, evaluated during the model's forward pass.

Arithmetic circuits and neural realizations of known computation. Arithmetic circuits represent polynomial computations as DAGs of additions and multiplications, with size and depth pro-

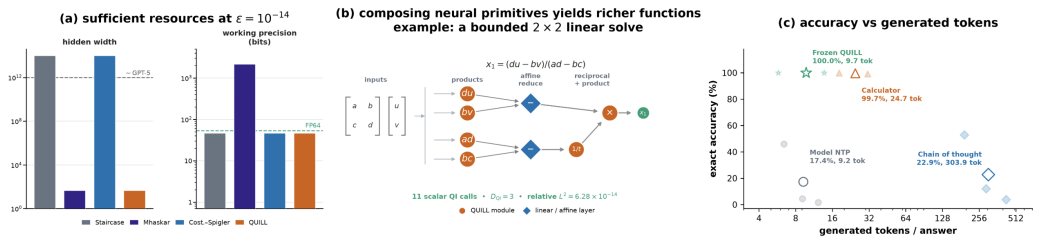


Figure 1: **QUILL makes high-precision computation realizable, compositional, and executable inside neural weights** (a) Resource requirements of explicit neural constructions at target error $\varepsilon = 10^{-14}$, comparing parameter count and per-parameter bit complexity. (b) Composition of neural primitives in a bounded 2×2 linear solve. QUILL multiplication and reciprocal modules are wired with affine operations to compute one solution coordinate, using 11 scalar QUILL calls, nonlinear depth 3, and achieving sampled relative L^2 error 6.28×10^{-14} in FP64.(c) Exact integer-squaring accuracy versus generated tokens for direct generation, chain-of-thought prompting, calculator use, and an installed QUILL module across Gemma-2-2B, Qwen2.5-7B, and Llama-3.1-8B.

viding standard measures of algebraic complexity (Shpilka & Yehudayoff, 2010). Prior MLP constructions have used explicit arithmetic modules compositionally: Yarotsky (2017) constructs approximate multiplication subnetworks as building blocks for smooth-function approximation. More directly, Arora et al. (2023) show that arithmetic circuits can be simulated by BaseConv gated-convolution models with controlled overhead in circuit depth and width. We use arithmetic circuits as a compilation method for composing high-precision MLP primitives into multivariate computations, with an end-to-end error bound for the resulting network (Theorem 4.1).

Constructive MLP approximation and quasi-interpolation. Universal approximation guarantees arbitrarily accurate MLPs without quantitative convergence rates (Cybenko, 1989), while constructive schemes establish explicit rates for sigmoidal networks (Jones, 1990). For analytic targets and target error ε, Mhaskar (1993; 1996) and De Ryck et al. (2021) obtain geometric or near-geometric convergence with width, but realizing these constructions requires working-precision bit complexity superlinear in $\log(1/\varepsilon)$ (Table 1). Alternatively, ChebNet (Tang et al., 2019) constructs stable neural realizations of Chebyshev polynomial approximations using deep ReQU networks, with depth growing logarithmically in the polynomial degree.

Classical quasi-interpolation constructs approximants from translated localized kernels (Buhmann & Jäger, 2022). Prior work adapts quasi-interpolation to shallow MLP operators (Anastassiou, 2011; Costarelli & Spigler, 2013; Chen & Cao, 2015). More specifically, Costarelli & Spigler (2015) construct sigmoidal approximants by approximating the target derivative with localized kernels and integrating the resulting expansion. QUILL combines these constructive approximation and quasi-interpolation ideas in the high-precision regime, obtaining geometric refinement in network width while remaining realizable in fixed-precision arithmetic down to a numerical floor (Theorem 3.1).

3 Constructing High-Precision MLPs

We develop QUILL using quasi-interpolation theory to construct high-precision MLPs (Section 3.1). Our main result, Theorem 3.1, proves that QUILL attains geometric convergence in width for bounded analytic targets while remaining realizable in floating-point arithmetic. For target error ε above the numerical floor, the required width scales as $O(\log(1/\varepsilon))$ (Corollary 3.1.1). The analysis also gives a rule for selecting the relative bandwidth—the product of feature spacing and pre-activation slope (Section 3.3). At fixed relative bandwidth, QUILL's slopes grow linearly with width. We show that small slopes suppress target-relevant kernel directions, slowing readout gradient descent (Theorem 3.2). Under sufficiently dispersed gradient signal across neurons, joint gradient descent fails to reach the construction's slope scale within the analyzed training budget (Theorem 3.3).

3.1 QUILL CONSTRUCTION

A quasi-interpolant viewpoint. Quasi-interpolation approximates functions using weighted sums of translated, localized kernels. QUILL connects this framework to MLPs through the derivatives of their activations. We present the tanh construction here; extensions to $GELU$ and $SiLU$ are developed in Appendix A.4. For a feature with center x_j and slope $\gamma > 0$,

$$\frac{d}{dx}\tanh\big(\gamma(x - x_j)\big) = \gamma\,\mathrm{sech}^2\big(\gamma(x - x_j)\big). \tag{1}$$

Approximating the target derivative with these kernels and integrating yields a tanh MLP.

Construction. We approximate a scalar target f on $\Omega = [-1, 1]$. Divide the interval into N uniform cells and extend the grid by R centers on each side:

$$h = \frac{2}{N}, \qquad x_j = -1 + jh, \qquad j = -R, \ldots, N + R. \tag{2}$$

The exterior centers provide *halo features* for correcting boundary effects. With a shared pre-activation slope $\gamma > 0$, define

$$q_{\mathbf{c}}(x) = b + \sum_{j=-R}^{N+R} w_j \tanh\big(\gamma(x - x_j)\big), \qquad \mathbf{c} = (b, w_{-R}, \ldots, w_{N+R}). \tag{3}$$

This is a width-$W = N + 2R + 1$ MLP, with prescribed input weight γ and hidden biases $-\gamma x_j$. Given $M = sN + 1$ equispaced observations $f(z_i)$ within Ω, recovering the readout is a linear least-squares problem; our implementation uses a reference-scaled SVD procedure (Appendix C):

$$\widehat{c} \in \arg\min_{c \in \mathbb{R}^{W+1}} \sum_{i=1}^{M} \left| q_c(z_i) - f(z_i) \right|^2. \tag{4}$$

Relative bandwidth. The dimensionless parameter $\lambda = \gamma h$ measures the relative overlap of neighboring features. Theorem 3.1 shows that λ governs an aliasing–conditioning tradeoff, which we balance through Equation 7.

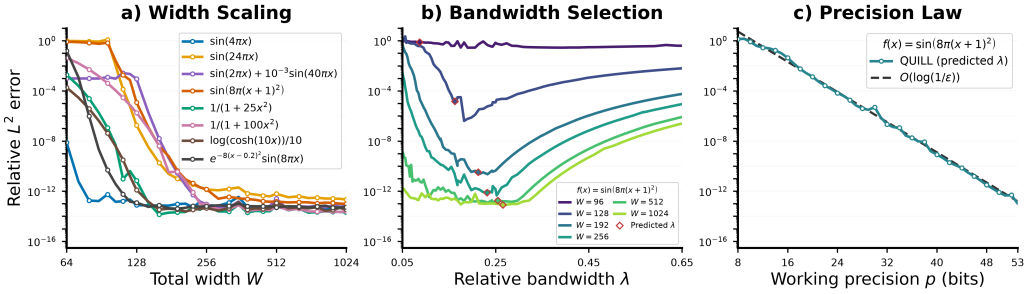


Figure 2: **QUILL achieves geometric refinement in a predictable bandwidth regime.** (a) Relative L_2 error versus width across target functions. QUILL converges geometrically with width until finite-precision recovery errors dominate. (b) Relative L_2 error versus relative bandwidth λ at several widths. The predicted bandwidth (red diamonds) (Equation 7) tracks the high-accuracy regime as the grid is refined. (c) Error versus working precision p, demonstrating logarithmic growth in the working precision needed to attain a target error. Appendix C documents calculations with custom solvers restricted to p-bit precision.

3.2 HIGH-PRECISION REMAINS NUMERICALLY REALIZABLE

The following theorem bounds the RMS error and readout magnitude of the constructed MLP. [1]

[1] We write $\|g\|_{2,\Omega} := \left(\frac{1}{2}\int_{-1}^{1} |g(x)|^2\,dx\right)^{1/2}$.

Construction	Maximum width	Largest parameter magnitude	Working precision (bits)	Hidden layers
Classical staircase (Jones, 1990)	$O(\varepsilon^{-1})$	$O(\varepsilon^{-1})$	$O(\log(1/\varepsilon))$	1
Mhaskar (1996)	$O(\log(1/\varepsilon))$	$\exp[\widetilde{O}(\log^2(1/\varepsilon))]$	$\widetilde{O}(\log^2(1/\varepsilon))$	1
Costarelli–Spigler (2015)	$O(\varepsilon^{-1})$	$O(\varepsilon^{-1})$	$O(\log(1/\varepsilon))$	1
ChebNet (2019)	$O(\log(1/\varepsilon))$	$O(1)$	$O(\log(1/\varepsilon))$	$O(\log\log(1/\varepsilon))$
QUILL	$O(\log(1/\varepsilon))$	$O(\log(1/\varepsilon))$	$O(\log(1/\varepsilon))$	1

Table 1: Resource bounds for explicit MLP constructions on a fixed univariate analytic class at error ε. Working precision covers construction and evaluation; parameter magnitudes include weights and biases. Construction details and derivations appear in Appendix D. QUILL's bounds apply in its admissible geometric-refinement regime. $\widetilde{O}$ hides powers of $\log\log(1/\varepsilon)$.

Theorem 3.1 (High-precision MLPs with stable readouts)**.** *Let $f : \Omega \to \mathbb{R}$ admit a bounded analytic extension to a complex neighborhood of Ω. Fix $0 < \lambda \le 1$, let $R = \lceil \sqrt{N} \rceil$, and choose a sufficiently large fixed observation density s. The computed network $\widehat{q}_W$, with readout $\overline{\mathbf{c}}_W$, satisfies*

$$\|f - \widehat{q}_W\|_{2,\Omega} \le \underbrace{C_f e^{-c_f(\lambda)W}}_{\textit{geometric refinement}} + \underbrace{\frac{C_f}{W} e^{-\pi^2/\lambda}}_{\textit{aliasing}} + \underbrace{\varepsilon_{\mathrm{num},W}}_{\textit{sampling and computation}}, \tag{5}$$

$$\|\overline{\mathbf{c}}_W\|_1 \le \underbrace{C_{\mathrm{out}}(f,\lambda)}_{\textit{width-independent readout bound}} + \underbrace{\varepsilon_{\mathrm{coeff},W}}_{\textit{coefficient formation}}. \tag{6}$$

The constants are independent of width. $\varepsilon_{\mathrm{num},W}$ bounds the contributions from sampling, readout recovery, coefficient formation, and floating-point evaluation, while $\varepsilon_{\mathrm{coeff},W}$ is the coefficient-formation allowance. The full proof and formal theorem statement are given in Appendix A.

Corollary 3.1.1 (Logarithmic-width high-precision MLPs)**.** *Under the conditions of Theorem 3.1, let ε be a target tolerance above the finite-precision floor. Choosing λ to make the aliasing contribution sufficiently small, QUILL attains*

$$\|f - \widehat{q}_W\|_{2,\Omega} \le \varepsilon, \qquad W = O\left(\log \frac{1}{\varepsilon}\right).$$

With aliasing controlled, numerical errors limit the accuracy gained from width refinement. Figure 2(a) shows geometric refinement followed by a numerical plateau, while Figure 2(c) shows that increasing working precision lowers the observed floor. Section 3.3 gives a practical rule for selecting λ from the sampled target and requested tolerance.

Table 1 compares QUILL's width, parameter magnitudes, and working-precision with previous constructions.

3.3 High-precision bandwidths are predictable

We predict the relative bandwidth λ from the sampled target and an effective aliasing tolerance ϵ_{eff} by balancing aliasing effects with readout amplification and boundary-correction cost

$$2\exp\left(-\frac{\pi^2}{\lambda_{\mathrm{pred}}}\right)\sinh\left(\frac{2\pi\bar{\omega}}{N\lambda_{\mathrm{pred}}}\right) = \epsilon_{\mathrm{eff}}, \tag{7}$$

where N is the number of interior grid cells and $\bar{\omega}$ is the amplitude-weighted mean absolute angular frequency, estimated from the sampled target's Fourier coefficients. Figure 2(b) shows that λ_{pred} tracks the observed high-precision bandwidth regime. Appendix B gives the derivation.

3.4 Small slopes delay readout learning

The high-precision relative bandwidth varies little with width (Figure 2(b)). The corresponding slopes $\gamma = \lambda/h$ therefore grow approximately linearly with width (Appendix B). Theorem 3.2 quantifies how smaller slopes delay readout gradient descent.

Theorem 3.2 (Slope-dependent output-error lower bound, informal)**.** *Fix the hidden geometry of Section 3.1 at common slope $\gamma > 0$. Readout GD from zero, with the prescribed kernel-scaled step size, satisfies*

$$\frac{\|f - q_n\|_{2,Z}^2}{\|f\|_{2,Z}^2} \geq \sum_i p_i(\gamma) \left(1 - \frac{\rho_i(\gamma)}{2}\right)^{2n}. \tag{8}$$

Here $p_i(\gamma)$ is the target-energy fraction in kernel direction i, $\rho_i(\gamma)$ is a slope-dependent upper bound on its normalized learning rate, derived from the kernel's spectral analysis. The formal statement, proof, and numerical validation appear in Appendix E.

Small slopes suppress target-relevant kernel directions: small $\rho_i(\gamma)$ with appreciable $p_i(\gamma)$ forces slow error decay. We next examine why joint training fails to acquire larger slopes.

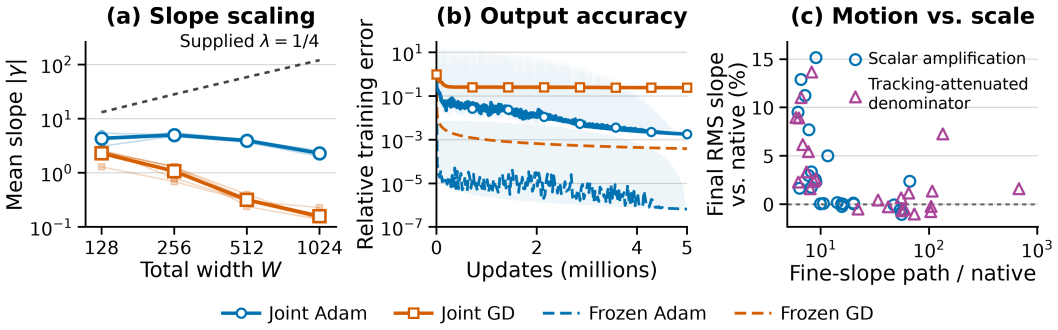


Figure 3: **Supplying QUILL geometry bypasses slow scale acquisition and improves optimization accuracy.** (a) Mean absolute pre-activation slopes learned by joint GD and Adam remain below the construction scale as width increases. (b) Readout-only training on a frozen dictionary with uniform centers and $\lambda = 0.25$ achieves lower relative training RMS error than joint training for both GD and Adam. Shading shows variation across the tested learning rates. (c) Adam interventions compare scalar amplification with attenuation of affine tracking in the second-moment denominator. The axes show accumulated nonlinear-fitting slope motion (horizontal, ratio to native Adam) and final RMS slope (vertical, percentage change from native Adam).

3.5 Why joint training acquires useful geometry slowly

Theorem 3.3 (Slow feature acquisition)**.** *Under the post-transient conditions of Appendix F, joint* tanh *GD on a nonzero sampled target satisfies, throughout the analyzed interval,*

$$\frac{\|f - q_n\|_{2,Z}}{\|f\|_{2,Z}} \geq \frac{[\|f_\perp\|_{2,Z} - Q_n(B_n)]_+}{\|f\|_{2,Z}}. \tag{9}$$

Here $f_\perp$ is the target component orthogonal to affine functions, B_n bounds the joint norm of slopes, hidden biases, and readouts, and $Q_n(B_n)$ bounds their non-affine output. The population RMS slope is at most $B_n/\sqrt{W}$.

After the affine fit, dispersed, moderate parameter energy weakens nonlinear sensitivity, while weak coupling of the leading cubic term to the target limits further growth. The resulting bound rules out high-precision fitting over the validated GD intervals (Appendix F). In Adam's higher-energy regime, affine-tracking updates can suppress nonlinear fitting through the shared second-moment denominator. Attenuating this contribution increases nonlinear-fitting motion without consistently increasing net slope growth (Figure 3(c); Appendix F.4).

QUILL bypasses this acquisition barrier by constructing the geometry and readout directly. Hence, we propose to use QUILL as a high-precision primitive for composing larger modules in language models.

4 COMPOSING QUILL PRIMITIVES INTO LARGE COMPUTATIONS

We now show how QUILL primitives can be assembled into larger numerical computations. We first develop a rigorous compilation theorem for polynomial arithmetic, then extend the same numerical compilation principle beyond polynomials, and finally consider discrete computation using explicitly constructed neural circuits.

4.1 COMPILING ARITHMETIC CIRCUITS

We begin with arithmetic circuits, which provide a standard representation for multivariate polynomial computation. An arithmetic circuit is a directed acyclic graph whose inputs are variables or constants and whose internal nodes, called *gates*, perform addition or multiplication. Intermediate values may be reused, and independent branches may be evaluated in parallel. Polynomial evaluation, matrix products, quadratic forms, and determinants all admit compact arithmetic circuits.

Circuit primitives. We use linear layers for addition, subtraction, and multiplication by constants, while the remaining non-affine multiplication gates are implemented from a QUILL x^2 module by polarization:

$$\widetilde{m}_\eta(x,y) = \widetilde{q}_\eta\left(\frac{x+y}{2}\right) - \widetilde{q}_\eta\left(\frac{x-y}{2}\right), \qquad |\widetilde{m}_\eta(x,y) - xy| \le 2\eta, \tag{10}$$

provided $\widetilde{q}_\eta$ uniformly approximates (Theorem A.3)t^2 to error η on the required bounded domain. Thus a QUILL x^2 module together with a linear layer suffices to compile any supplied polynomial arithmetic circuit.

Theorem 4.1 (Compilation of bounded arithmetic circuits)**.** *Let Φ be a supplied arithmetic circuit. Suppose addition, subtraction, and multiplication by constants are implemented exactly with linear layers, while each multiplication gate is replaced by a QUILL multiplier. Let $S_\times(\Phi)$ denote its number of non-affine multiplication gates, $D_\times(\Phi)$ its multiplicative depth, and κ_Φ the circuit-dependent amplification of local multiplication error defined by the gatewise recurrence in Appendix G. Suppose the exact and compiled intermediate values remain within prescribed bounded domains.*

Then, for any target error ε for which the required local accuracy $\delta = O(\varepsilon/\kappa_\Phi)$ remains above the QUILL multiplier's finite-precision floor, there exists a QUILL-compiled network $N_{\Phi,\varepsilon}$ satisfying

$$\|N_{\Phi,\varepsilon} - f_\Phi\|_{\infty,\Omega} \le \varepsilon.$$

Thus, using the geometric convergence established in Theorem 3.1,

$$S_{\mathrm{hid}} = O\Big(S_\times(\Phi)\left[1 + \log_+\Big(\frac{\kappa_\Phi}{\varepsilon}\Big)\right]\Big), \qquad D_{\mathrm{nl}} = O(D_\times(\Phi)),$$

where S_{hid} is the total number of constructed hidden units and D_{nl} the number of sequential nonlinear layers.

Thus the compiled network inherits the structure of the source circuit: $S_\times(\Phi)$ governs the number of multiplication modules, $D_\times(\Phi)$ their sequential depth, and κ_Φ the error propagation. Because each QUILL multiplier requires only logarithmic width in its target accuracy, error amplification increases network size only logarithmically in κ_Φ/ε. Full details are given in Appendix G.

Polynomials. Because arithmetic circuits exactly represent multivariate polynomials, Theorem 4.1 reduces their compiled-network complexity to the multiplication count $S_\times$ and multiplicative depth $D_\times$ of the chosen circuit. For a polynomial with M monomials and total degree at most d, a direct balanced construction of each monomial gives

$$S_\times = O(Md), \qquad D_\times = O(\log d),$$

with smaller circuits possible when intermediate quantities are reused or when common subcomputations are implemented directly by QUILL modules.

We demonstrate this compilation procedure on structured polynomial computations, including dot products, squared distances, quadratic forms, tensor contractions, and determinants. A 3×3 determinant illustrates the resulting circuit structure: with $S_\times = 12$ multiplication gates and multiplicative depth $D_\times = 2$, the resulting compiled network achieves relative L_2 error 1.55×10^{-14}. Across all tested polynomial programs, the largest sampled relative L_2 error is 6.46×10^{-14} in FP64.

Table 2: **Representative compiled numerical and discrete programs.** Numerical rows report held-out FP64 approximation or implementation error; discrete rows report decoded accuracy.

Class	Program	Constructed primitives	Result
Scalar	Normal CDF on $[-3, 3]$	QI CDF	rel. L^2: 2.09×10^{-15}
Polynomial	3×3 determinant	QI square $\rightarrow$ multiply	rel. L^2: 1.55×10^{-14}
Non-polynomial	3D Gaussian RBF	QI square, exp	rel. L^2: 5.28×10^{-14}
Algorithm	Newton reciprocal, 5 steps	QI square $\rightarrow$ multiply	impl. max: 3.11×10^{-14}
Discrete	Eight-bit rule evaluation	tanh threshold gates	100% decode
Discrete	Triangle detection, 5 vertices	tanh threshold gates	100% decode

4.2 Beyond polynomial arithmetic

Expanding the primitive library. Polynomial circuits require only addition and multiplication. Expanding the primitive library with a small number of additional scalar functions greatly enlarges the class of computations that can be represented. QUILL can approximate bounded primitives such as $1/x$, $\sin(x)$, and $1/\sqrt{x}$, which can be inserted as additional gates while retaining the same compositional error analysis.

This is enough to express common non-polynomial operations such as Gaussian radial basis functions, softmax, cosine similarity, and bounded matrix inversion. We demonstrate this by compiling these functions directly from the expanded primitive library; 3D Gaussian RBF and three-logit softmax achieve relative L_2 errors of 5.28×10^{-14} and 4.36×10^{-14}, respectively, in FP64.

Fixed-step numerical algorithms. The same framework extends from closed-form computations to numerical algorithms with fixed iteration numbers. Unrolling the iterations produces a finite feed-forward computation graph whose update equations can be compiled from the same primitive library.

For Newton reciprocal iteration,

$$r_{k+1} = r_k(2 - xr_k),$$

unrolling K steps yields a finite graph with two dependent multiplications per iteration. Four compiled steps match the exact four-step iteration to 3.02×10^{-14} implementation error. With five steps, the compiled program has implementation error 3.11×10^{-14} relative to the exact five-step iteration, while its total error relative to $1/x$ is 3.06×10^{-14}. We similarly compile an eight-node quadrature rule, matching its prescribed finite computation to 4.42×10^{-14}.

4.3 Discrete computation and Boolean circuits

The same compile-into-weights perspective also extends to discrete computation. Using explicitly constructed tanh MLP gates, we can realize supplied Boolean circuits directly in neural-network weights. In particular, any polynomial-size, constant-depth AC^0 circuit compiles into a polynomial-size, constant-depth MLP with exact correctness after Boolean decoding. The construction and proof are given in Appendix G.1. We also demonstrate larger finite discrete programs built from these primitives. Random Boolean circuits decode correctly through depth 32, bitonic sorting exactly sorts up to 32 eight-bit integers, and a min-plus shortest-path construction returns correct rounded distances and valid predecessors through $n = 16$. These latter examples are finite-program demonstrations rather than asymptotic AC^0 claims.

5 Proof-of-Concept LLM Installation

Finally, as a proof of concept, we attach a frozen width-80 QUILL x^2 module to Gemma-2-2B-IT, Qwen2.5-7B-Instruct, and Llama-3.1-8B-Instruct and invoke it on a supplied operand during inference. We choose x^2 for this demonstration. Appendix A.4 studies other unary targets and constructions across other activation functions.

Table 3: **Accuracy @ mean generated output tokens for integer squaring.** QUILL performs the numerical operation internally. For CoT, we report accuracy on completed generations only; 187/2250 cases reached the 1,024-token limit.

Digits	Model	Direct	CoT	Calculator	QUILL
3	Gemma-2-2B-IT	35.2% @ 9.1	44.4% @ 104.9	100% @ 22.8	100% @ 5.8
5	Gemma-2-2B-IT	4.0% @ 11.7	2.8% @ 150.8	100% @ 30.7	100% @ 9.7
7	Gemma-2-2B-IT	0.0% @ 16.4	1.2% @ 456.7	97.6% @ 38.7	100% @ 13.7
3	Qwen2.5-7B-Instruct	79.2% @ 6.7	79.2% @ 291.4	100% @ 19.8	100% @ 5.8
5	Qwen2.5-7B-Instruct	7.2% @ 10.2	28.4% @ 366.3	100% @ 27.7	100% @ 9.7
7	Qwen2.5-7B-Instruct	4.8% @ 12.7	9.6% @ 239.5	100% @ 35.7	100% @ 13.7
3	Llama-3.1-8B-Instruct	23.6% @ 3.4	35.2% @ 184.0	99.6% @ 12.0	100% @ 5.8
5	Llama-3.1-8B-Instruct	2.0% @ 5.5	4.8% @ 359.1	100% @ 15.7	100% @ 9.7
7	Llama-3.1-8B-Instruct	0.4% @ 7.6	0.4% @ 582.9	100% @ 19.0	100% @ 13.7
All	**All models**	**17.4% @ 9.2**	**20.7% @ 303.9**	**99.7% @ 24.7**	**100% @ 9.7**

5.1 Installation

Transformer models contain MLP blocks, making them a natural setting for inserting a constructed MLP module. Our pathway has three stages. First, an engineered reader reconstructs the supplied operand from its layer-0 representation. Second, the QUILL module evaluates x^2 during the layer-20 forward call. Third, a serializer rounds the result and writes the corresponding output representation to the final hidden state, where the frozen language-model head decodes it into answer tokens.

Triggering, operand extraction, and serialization are fixed by the implementation; the numerical function itself is evaluated by the constructed QUILL weights. The host transformer runs in BF16, and the QUILL module runs in FP64.

5.2 Results

We evaluate x^2 for 3-, 5-, and 7-digit integers in Gemma-2-2B-IT, Qwen2.5-7B-Instruct, and Llama-3.1-8B-Instruct. We compare four methods: direct generation, chain-of-thought prompting, calculator use, and the QUILL pathway.

Across all nine model–digit settings, the QUILL pathway achieves 100% exact accuracy with 9.7 generated answer tokens on average. Calculator use achieves 99.7% exact accuracy with 24.7 generated tokens. Because $187/2{,}250$ chain-of-thought generations hit our 1,024-token limit, the reported mean token count understates the true average and should be interpreted as a lower bound. Together, these results provide a proof of concept for in-weight execution inside a pretrained language model.

6 Discussion

QUILL establishes a constructive route to numerically realizable high-precision computation in ordinary MLP weights. For bounded analytic functions, QUILL achieves geometric refinement toward a finite-precision floor. QUILL also provides a reference for understanding training: gradient descent acquires the hidden-layer scales supporting high precision too slowly, while direct construction bypasses this acquisition stage. QUILL primitives compose into in-weight programs ranging from determinants and softmax to fixed-step numerical methods and bounded graph algorithms. For bounded arithmetic computations, we establish end-to-end accuracy and resource bounds, extending high-precision neural construction from scalar functions to entire numerical programs. The proof-of-concept installation demonstrates that a constructed numerical function can execute inside a frozen pretrained language model, returning exact squares on every evaluated input. This provides a route to adding modular numerical capabilities *in-weight* in modern language models. Importantly, the current proof-of-concept pathway uses engineered routing and input/output handling. Future work should explore learning to invoke and compose these modules inside LLMs. Together, these results establish numerical construction as a means of supplying reusable computational subroutines alongside learned components within neural architectures.

AI USE STATEMENT

In this work, we used generative AI tools for plotting, literature search, spellchecking, brainstorming, and independent proof checking. We have not used generative AI tools to build core ideas, fundamental proof sketches, or theoretical breakthroughs, and remaining disclosures are not applicable to this work. We have reviewed all AI-assisted work by scrutinizing proofs, consistently checked collaborator's code, and rewritten ideas in our own words. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This work involves no human subjects or personal data; all experiments use synthetic inputs. The language-model experiments use publicly released open-weight models under their respective licenses, and the models remain frozen. Because QUILL writes hand-constructed computations into network weights, it could in principle be used to add hard-to-detect behavior to a model, but our installed module is an explicitly attached component that computes only x^2, and we see no risks beyond those of existing model-editing methods. We are not aware of any conflicts of interest.

REPRODUCIBILITY STATEMENT

Formal statements and proofs of Theorems 3.1, 3.2, 3.3, and 4.1 appear in Appendices A, E, F, and G, respectively. Appendix B derives the bandwidth rule, Appendix C specifies the precision-controlled arithmetic and solver, and Appendix D derives the comparison bounds in Table 1. Experimental settings are described in the figure captions and appendices. Source code is included in the supplementary material.

ACKNOWLEDGMENTS

Anonymous during submission.

REFERENCES

Ben Adcock and Daan Huybrechs. Frames and numerical approximation II: Generalized sampling. *Journal of Fourier Analysis and Applications*, 26(6):87, 2020. doi: 10.1007/s00041-020-09796-w. URL https://doi.org/10.1007/s00041-020-09796-w.

George A. Anastassiou. Univariate hyperbolic tangent neural network approximation. *Mathematical and Computer Modelling*, 53(5):1111–1132, 2011. ISSN 0895-7177. doi: https://doi.org/10.1016/j.mcm.2010.11.072. URL https://www.sciencedirect.com/science/article/pii/S0895717710005649.

Edward Anderson, Zhaojun Bai, Christian Bischof, L. Susan Blackford, James Demmel, Jack Dongarra, Jeremy Du Croz, Anne Greenbaum, Sven Hammarling, Alan McKenney, and Danny Sorensen. *LAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 3 edition, 1999. doi: 10.1137/1.9780898719604. URL https://doi.org/10.1137/1.9780898719604.

Simran Arora, Sabri Eyuboglu, Aman Timalsina, Isys Johnson, Michael Poli, James Zou, Atri Rudra, and Christopher Ré. Zoology: Measuring and improving recall in efficient language models, 2023. URL https://arxiv.org/abs/2312.04927.

Alex Bie, Travis Dick, Alex Kulesza, Prabhakar Raghavan, Vinod Raman, and Sergei Vassilvitskii. Ai-rithmetic, 2026. URL https://arxiv.org/abs/2602.10416.

Claude Brezinski and Michela Redivo-Zaglia. New representations of padé, padé-type, and partial padé approximants. *Journal of Computational and Applied Mathematics*, 284:69–77, 2015. doi: 10.1016/j.cam.2014.07.007. URL https://doi.org/10.1016/j.cam.2014.07.007.

Martin Buhmann and Janin Jäger. *Quasi-Interpolation*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2022.

Claudio Canuto, M. Yousuff Hussaini, Alfio Quarteroni, and Thomas A. Zang. *Spectral Methods: Fundamentals in Single Domains*. Scientific Computation. Springer, Berlin, Heidelberg, 2006. ISBN 978-3-540-30725-9. doi: 10.1007/978-3-540-30726-6. URL https://doi.org/10.1007/978-3-540-30726-6.

Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks, 2023. URL https://arxiv.org/abs/2211.12588.

Zhi-xiang Chen and Fei-long Cao. The construction and approximation of feedforward neural network with hyperbolic tangent function. *Applied Mathematics-A Journal of Chinese Universities*, 30(2):151–162, 2015. doi: 10.1007/s11766-015-3000-9. URL https://doi.org/10.1007/s11766-015-3000-9.

Jeremy M. Cohen, Behrooz Ghorbani, Shankar Krishnan, Naman Agarwal, Sourabh Medapati, Michal Badura, Daniel Suo, David Cardoze, Zachary Nado, George E. Dahl, and Justin Gilmer. Adaptive gradient methods at the edge of stability, 2022. URL https://arxiv.org/abs/2207.14484v2. Version 2, revised April 2024.

D. Costarelli and R. Spigler. Constructive approximation by superposition of sigmoidal functions. *Analysis in Theory and Applications*, 29(2):169–196, 2013. doi: 10.4208/ata.2013.v29.n2.8. URL https://doi.org/10.4208/ata.2013.v29.n2.8.

Danilo Costarelli and Renato Spigler. Approximation by series of sigmoidal functions with applications to neural networks. *Annali di Matematica Pura ed Applicata (1923 -)*, 194(1):289–306, 2015. doi: 10.1007/s10231-013-0378-y. URL https://doi.org/10.1007/s10231-013-0378-y.

George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989. doi: 10.1007/BF02551274.

Tim De Ryck, Samuel Lanthaler, and Siddhartha Mishra. On the approximation of functions by tanh neural networks. *Neural Networks*, 143:732–750, November 2021. ISSN 0893-6080. doi: 10.1016/j.neunet.2021.08.015. URL http://dx.doi.org/10.1016/j.neunet.2021.08.015.

Florian Dietz and Dietrich Klakow. Igc: Integrating a gated calculator into an llm to solve arithmetic tasks reliably and efficiently, 2025. URL https://arxiv.org/abs/2501.00684.

Owen Dugan, Donato Manuel Jimenez Beneto, Charlotte Loh, Zhuo Chen, Rumen Dangovski, and Marin Soljačić. Occamllm: Fast and exact language model arithmetic in a single step, 2024. URL https://arxiv.org/abs/2406.06576.

Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. Pal: Program-aided language models, 2023. URL https://arxiv.org/abs/2211.10435.

Walter Gautschi. On inverses of vandermonde and confluent vandermonde matrices. *Numerische Mathematik*, 4:117–123, 1962. doi: 10.1007/BF01386302. URL https://doi.org/10.1007/BF01386302.

Per Christian Hansen. The truncated SVD as a method for regularization. *BIT Numerical Mathematics*, 27(4):534–553, 1987. doi: 10.1007/BF01937276. URL https://doi.org/10.1007/BF01937276.

Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2 edition, 2002. doi: 10.1137/1.9780898718027. URL https://doi.org/10.1137/1.9780898718027.

Mohsin Javed and Lloyd N. Trefethen. A trapezoidal rule error bound unifying the euler–maclaurin formula and geometric convergence for periodic functions. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 470(2161):20130571, 2014. doi: 10.1098/rspa.2013.0571. URL https://doi.org/10.1098/rspa.2013.0571.

L.K. Jones. Constructive approximations for neural networks by sigmoidal functions. *Proceedings of the IEEE*, 78(10):1586–1589, 1990. doi: 10.1109/5.58342.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. URL https://arxiv.org/abs/1412.6980.

Thang Luong, Dawsen Hwang, Hoang H. Nguyen, Golnaz Ghiasi, Yuri Chervonyi, Insuk Seo, Junsu Kim, Garrett Bingham, Jonathan Lee, Swaroop Mishra, Alex Zhai, Clara Huiyi Hu, Henryk Michalewski, Jimin Kim, Jeonghyun Ahn, Junhwi Bae, Xingyou Song, Trieu H. Trinh, Quoc V. Le, and Junehyuk Jung. Towards robust mathematical reasoning, 2025. URL https://arxiv.org/abs/2511.01846.

H. N. Mhaskar. Neural networks for optimal approximation of smooth and analytic functions. *Neural Computation*, 8(1):164–177, 1996. doi: 10.1162/neco.1996.8.1.164.

Hrushikesh Narhar Mhaskar. Approximation properties of a multilayered feedforward artificial neural network. *Advances in Computational Mathematics*, 1(1):61–80, 1993.

Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools, 2023. URL https://arxiv.org/abs/2302.04761.

Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Found. Trends Theor. Comput. Sci.*, 5:207–388, 2010. URL https://api.semanticscholar.org/CorpusID:86398017.

Shanshan Tang, Bo Li, and Haijun Yu. ChebNet: Efficient and stable constructions of deep neural networks with rectified power units via chebyshev approximations, 2019. URL https://arxiv.org/abs/1911.05467.

Lloyd N. Trefethen. *Approximation Theory and Approximation Practice, Extended Edition*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2019. doi: 10.1137/1.9781611975949. URL https://epubs.siam.org/doi/abs/10.1137/1.9781611975949.

Lloyd N. Trefethen and David Bau. *Numerical Linear Algebra, Twenty-fifth Anniversary Edition*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2022. doi: 10.1137/1.9781611977165. URL https://epubs.siam.org/doi/abs/10.1137/1.9781611977165.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023. URL https://arxiv.org/abs/2201.11903.

Dmitry Yarotsky. Error bounds for approximations with deep relu networks, 2017. URL https://arxiv.org/abs/1610.01145.

A CONSTRUCTION THEORY

We prove Theorem A.1, the precise statement underlying Theorem 3.1. We first construct an accurate reference network in the prescribed tanh dictionary, and then establish its recovery from in-domain samples and the resulting continuous-RMS bound. Section A.4 extends the representation argument to exact GELU and SiLU.

A.1 SETUP AND THEOREM STATEMENT

Let $\Omega = [-1, 1]$, $I_\Delta = [-1 - \Delta, 1 + \Delta]$, and $\delta, \Delta, B > 0$. The class $\mathcal{A}_{\delta,B}(I_\Delta)$ consists of the functions real on I_Δ, holomorphic on

$$U_\delta(I_\Delta) = \{z \in \mathbb{C} : \operatorname{dist}(z, I_\Delta) < \delta\},$$

and bounded there by B. We use

$$\|g\|_{2,\Omega}^2 = \frac{1}{2}\int_{-1}^{1} |g(x)|^2\, dx, \qquad \|g\|_\infty = \sup_{x\in\Omega} |g(x)|.$$

Fix $0 < \lambda \le 1$ and $a_H > 0$, and define

$$h = \frac{2}{N}, \quad x_j = -1 + jh, \quad R = \lceil a_H\sqrt{N}\rceil, \quad \gamma = \frac{\lambda}{h}, \quad W = N + 2R + 1. \tag{A.1}$$

The tanh dictionary is

$$q_c(x) = b_c + \sum_{j=-R}^{N+R} w_j \tanh(\gamma(x - x_j)), \qquad \|c\|_1 = |b_c| + \sum_j |w_j|. \tag{A.2}$$

Call this geometry admissible if

$$(R + \tfrac{1}{2})h \le \Delta, \quad \gamma\delta \ge 4\pi, \quad \lambda R \ge 2\log 2, \quad N \ge 16\max\{1, a_H^2\}. \tag{A.3}$$

At observation density $s \in \mathbb{N}$, use only core samples

$$z_i = -1 + \frac{ih}{s}, \quad 0 \le i \le Ns, \qquad M = Ns + 1, \qquad \|g\|_{2,X}^2 = \frac{1}{M}\sum_{i=0}^{M-1} |g(z_i)|^2. \tag{A.4}$$

There are $n = W + 1$ readout coordinates, including the output bias. The sampled algorithm is reference-scaled truncated SVD, as specified in Section A.3; it is not an unregularized least-squares minimizer. It uses real target samples and deterministic coefficient envelopes; the reference construction's high derivatives and complex target values are not evaluated by the sampled algorithm. The guarantees below concern successful runs satisfying the stated numerical conditions.

Theorem A.1 (Approximation and finite-precision recovery)**.** *Let $f \in \mathcal{A}_{\delta,B}(I_\Delta)$ and let (A.1) be admissible. Suppose the numerical and sampling conditions in Section A.3 and the input-rounding conditions in Section A.3.5 hold, with $M \ge n$. Every returned network satisfies*

$$\|\widehat{c}_W\|_1 \le 2BK_{\mathrm{ref}} + \epsilon_{\mathrm{coeff,W}}, \tag{A.5}$$

$$\|f - \widehat{q}_W\|_{2,\Omega} \le C_f\left(e^{-c_{\mathrm{app}}W} + \frac{e^{-\pi^2/\lambda}}{W}\right) + \epsilon_{\mathrm{rec,W}}, \tag{A.6}$$

$$c_{\mathrm{app}} = \tfrac{1}{8}\min\{\pi\delta, \lambda a_H^2\}, \qquad C_f = \frac{512\pi B}{\lambda\delta}. \tag{A.7}$$

Here K_{ref} is a width-independent bound on the sum of the reference envelopes in (A.54). The nonnegative allowances $\epsilon_{\mathrm{coeff,W}}$ and $\epsilon_{\mathrm{rec,W}}$ are defined in (A.84). All input-rounding contributions are included in these allowances; in particular, C_f does not depend on W.

A.2 TANH REFERENCE NETWORK

Local analytic representation. Set

$$d = \frac{\pi}{2\gamma}, \quad r = (R + \tfrac{1}{2})h, \quad J = [a, b] = [-1 - r, 1 + r], \quad x_0 = -1. \tag{A.8}$$

The W centers in (A.1) are the midpoints of the width-h cells partitioning J, and $Wh = 2 + 2r$. Define

$$K_\gamma(t) = \frac{\gamma}{2}\operatorname{sech}^2(\gamma t), \qquad a_\gamma(z) = \frac{f(z + id) - f(z - id)}{2id} = \frac{1}{2d}\int_{-d}^{d} f'(z + iv)\, dv. \tag{A.9}$$

Approximating a derivative by localized kernels and integrating the expansion is a standard route to sigmoidal approximation Costarelli & Spigler (2015); Buhmann & Jäger (2022). Here the shifted-value formula gives a density adapted to a finite analytic neighborhood.

Lemma A.1.1 (Local density bound)**.** *If $\sigma + d < \delta$, the density is holomorphic near $J + i[-\sigma, \sigma]$, is real on J, and satisfies*

$$\sup_{t \in J,\ |y| \le \sigma} |a_\gamma(t + iy)| \le M_\sigma := \frac{B}{\delta - \sigma - d}. \tag{A.10}$$

Proof. Cauchy's estimate gives $|f'(t + i(y + v))| \le B/(\delta - \sigma - d)$ for $|v| \le d$; average in (A.9). The shifted formula gives holomorphy, and $f(\bar{z}) = \overline{f(z)}$ gives reality on J. □

For $x \in \Omega$, put

$$D_x(z) = \tfrac{1}{2}[\tanh(\gamma(x - z)) - \tanh(\gamma(x_0 - z))], \qquad F_x(z) = a_\gamma(z) D_x(z), \tag{A.11}$$

$$(\mathcal{C}_h f)(x) = f(x_0) + \int_a^b F_x(t)\, dt, \tag{A.12}$$

$$H_x(z) = \coth(\gamma(x - z)) - \coth(\gamma(x_0 - z)), \tag{A.13}$$

$$(\mathcal{B}_h f)(x) = \frac{1}{4d}\int_{-d}^{d} [f(a + iy) H_x(a + iy) - f(b + iy) H_x(b + iy)]\, dy. \tag{A.14}$$

Lemma A.1.2 (Finite local representation)**.** *For $x \in \Omega$,*

$$f(x) = \mathcal{C}_h f(x) + \mathcal{B}_h f(x), \qquad \|\mathcal{B}_h f\|_\infty \le B[\coth(\gamma r) - \coth(\gamma(r + 2))] \le \frac{2B}{e^{2\gamma r} - 1}. \tag{A.15}$$

Proof. The assertion at $x = x_0$ is immediate. Otherwise, $\gamma d = \pi/2$ implies $H_x(t \pm id) = 2D_x(t)$, and hence

$$\int_a^b F_x(t)\, dt = \frac{1}{4id}\left[\int_a^b f(t + id) H_x(t + id)\, dt - \int_a^b f(t - id) H_x(t - id)\, dt\right].$$

On the positively oriented rectangle $J + i[-d, d]$, the residues of $f H_x$ at x and x_0 are $-f(x)/\gamma$ and $f(x_0)/\gamma$. Separating the two horizontal integrals and the vertical sides gives (A.15), since $4d\gamma = 2\pi$. For $\ell = x - x_0 \in [0, 2]$, the expansion $\coth w = 1 + 2\sum_{k \ge 1} e^{-2kw}$ for $\Re w > 0$ and oddness give

$$|H_x(a + iy)| \le \coth(\gamma r) - \coth(\gamma(r + \ell)),$$
$$|H_x(b + iy)| \le \coth(\gamma(r + 2 - \ell)) - \coth(\gamma(r + 2)).$$

Each is at most $\coth(\gamma r) - \coth(\gamma(r + 2))$. Insert these bounds in (A.14) and use $\coth t - 1 = 2/(e^{2t} - 1)$. □

We correct $\mathcal{B}_h f$ together with the endpoint contribution introduced by discretization.

Midpoint discretization. We use the finite-interval characteristic-function argument of (Javed & Trefethen, 2014, Secs. 4–5). The following formulation fixes the orientations and retains the residues of a meromorphic integrand. Let $b = a + m_q h$, $t_j = a + (j + \frac{1}{2})h$, and define

$$Q_h F = h \sum_{j=0}^{m_q - 1} F(t_j), \qquad IF = \int_a^b F(t)\, dt, \qquad \zeta_0 = a + h/2.$$

For F meromorphic near $[a, b] + i[-\sigma, \sigma]$, with finitely many poles and none on the boundary or real segment, put

$$\mathcal{K}_h(z) = \frac{1}{2i} \cot \frac{\pi(z - \zeta_0)}{h}, \qquad R_h^{\pm} = \mathcal{K}_h \pm \tfrac{1}{2}. \tag{A.16}$$

Let $\mathcal{P}_{\pm}$ be its upper/lower pole sets, and set

$$\mathcal{H}_h[F] = \int_a^b [R_h^-(t - i\sigma)F(t - i\sigma) - R_h^+(t + i\sigma)F(t + i\sigma)]\, dt, \tag{A.17}$$

$$\mathcal{V}_h[F] = i \int_0^{\sigma} \frac{F(b + iy) - F(a + iy) - F(b - iy) + F(a - iy)}{1 + e^{2\pi y/h}}\, dy, \tag{A.18}$$

$$\mathcal{P}_h[F] = 2\pi i \left[\sum_{p \in \mathcal{P}_+} \operatorname{Res}(R_h^+ F; p) + \sum_{p \in \mathcal{P}_-} \operatorname{Res}(R_h^- F; p) \right]. \tag{A.19}$$

For a simple pole, $\operatorname{Res}(R_h^{\pm} F; p) = R_h^{\pm}(p) \operatorname{Res}(F; p)$.

Lemma A.1.3 (Signed midpoint identity)**.**

$$Q_h F - IF = \mathcal{H}_h[F] + \mathcal{V}_h[F] - \mathcal{P}_h[F], \tag{A.20}$$

and, in their respective half-planes,

$$|R_h^+(z)| \le \frac{1}{e^{2\pi \Im z/h} - 1}, \qquad |R_h^-(z)| \le \frac{1}{e^{2\pi |\Im z|/h} - 1}. \tag{A.21}$$

Proof. Apply the contour argument in (Javed & Trefethen, 2014, Secs. 4–5) with (A.16). Equivalently, the positively oriented upper path $b \to b + i\sigma \to a + i\sigma \to a$ and lower path $a \to a - i\sigma \to b - i\sigma \to b$ give

$$\int_{\Gamma_+} F = 2\pi i \sum_{\mathcal{P}_+} \operatorname{Res} F - IF, \qquad \int_{\Gamma_-} F = 2\pi i \sum_{\mathcal{P}_-} \operatorname{Res} F + IF.$$

Every grid pole of $\mathcal{K}_h$ has residue $h/(2\pi i)$. Apply the residue theorem to $\mathcal{K}_h F$ on the combined contour and substitute $\mathcal{K}_h = R_h^+ - 1/2$ above and $R_h^- + 1/2$ below. At the half-cell endpoints, $R_h^+(a + iy) = R_h^+(b + iy) = (1 + e^{2\pi y/h})^{-1}$ and the lower values are their negatives. These identities give (A.20). Writing $w = e^{2\pi i(z - \zeta_0)/h}$ gives $R_h^+ = w/(w - 1)$ and $R_h^- = 1/(w - 1)$, proving (A.21). □

For tanh choose

$$k = \left\lceil \frac{\delta}{8d} \right\rceil, \qquad \sigma = 2kd. \tag{A.22}$$

Admissibility implies $d \le \delta/8$ and

$$\delta/4 \le \sigma \le \delta/2, \quad \sigma + d \le 5\delta/8, \quad M_\sigma \le 4B/\delta. \tag{A.23}$$

Also $\gamma\sigma = k\pi$, so $D_x(t \pm i\sigma) = D_x(t)$.

Lemma A.1.4 (Tanh quadrature bounds)**.** *Uniformly for $x \in \Omega$,*

$$|\mathcal{H}_h[F_x]| \le \frac{4M_\sigma}{e^{2\pi\sigma/h} - 1}, \tag{A.24}$$

$$|\mathcal{P}_h[F_x]| \le \frac{4\pi M_\sigma}{\gamma} \sum_{\ell=0}^{k-1} \frac{1}{e^{(2\ell+1)\pi^2/\lambda} - 1}, \tag{A.25}$$

$$|\mathcal{V}_h[F_x]| \le \frac{4M_\sigma h \log 2}{\pi(e^{2\gamma r} - 1)}. \tag{A.26}$$

Proof. On the real axis, $D_x(t) = \int_{x_0}^{x} K_\gamma(v-t)\,dv \geq 0$ and $\int_{\mathbb{R}} D_x(t)\,dt = x - x_0 \leq 2$. Together with imaginary periodicity and (A.21), this proves (A.24). The pole families are

$$x \pm i(2\ell+1)d, \qquad x_0 \pm i(2\ell+1)d, \qquad 0 \leq \ell < k.$$

Their residues in D_x have modulus $1/(2\gamma)$, and $2\pi(2\ell+1)d/h = (2\ell+1)\pi^2/\lambda$, giving (A.25). At the vertical sides, subtract the common saturated value of the two tanhs and use $\tanh w - 1 = -2/(e^{2w}+1)$ for $\Re w > 0$, together with oddness. This gives $|D_x(a \pm iy)|, |D_x(b \pm iy)| \leq 2/(e^{2\gamma r}-1)$. Now use (A.18) and $\int_0^\infty (1+e^{2\pi y/h})^{-1}\,dy = h\log 2/(2\pi)$. □

The midpoint network is

$$q_h(x) = f(x_0) + h\sum_{j=-R}^{N+R} a_\gamma(x_j) D_x(x_j). \tag{A.27}$$

Its weights $w_j^{\mathrm{base}} = h a_\gamma(x_j)/2$ satisfy

$$|w_j^{\mathrm{base}}| \leq \frac{Bh}{\delta}, \qquad \sum_j |w_j^{\mathrm{base}}| \leq \frac{2B(1+\Delta)}{\delta}. \tag{A.28}$$

Subtracting (A.20) from the finite representation gives

$$f - q_h = \mathcal{B}_h^{\mathrm{eff}} f + \mathcal{P}_h[F_x] - \mathcal{H}_h[F_x], \qquad \mathcal{B}_h^{\mathrm{eff}} f := \mathcal{B}_h f - \mathcal{V}_h[F_x]. \tag{A.29}$$

Boundary correction using the existing halo. Set

$$z_-(x) = e^{-2\gamma(x-a)}, \quad z_+(x) = e^{-2\gamma(b-x)}, \quad \theta = e^{-\gamma r}, \quad z_{\pm,0} = z_\pm(x_0). \tag{A.30}$$

All these arguments lie in $(0,\theta^2]$. For $\ell \geq 1$, define

$$\mu_\ell^- = \frac{1}{2d}\int_{-d}^{d} f(a+iy)e^{2i\ell\gamma y}\,dy, \qquad \mu_\ell^+ = \frac{1}{2d}\int_{-d}^{d} f(b+iy)e^{-2i\ell\gamma y}\,dy, \tag{A.31}$$

$$\nu_\ell^- = i(-1)^\ell \int_0^\sigma \frac{a_\gamma(a-iy)e^{-2i\ell\gamma y} - a_\gamma(a+iy)e^{2i\ell\gamma y}}{1+e^{2\pi y/h}}\,dy, \tag{A.32}$$

$$\nu_\ell^+ = i(-1)^\ell \int_0^\sigma \frac{a_\gamma(b-iy)e^{2i\ell\gamma y} - a_\gamma(b+iy)e^{-2i\ell\gamma y}}{1+e^{2\pi y/h}}\,dy. \tag{A.33}$$

Write $g_\ell^\pm = \mu_\ell^\pm - \nu_\ell^\pm$, $G^\pm(z) = \sum_{\ell\geq 1} g_\ell^\pm z^\ell$, and $D_\lambda = \pi/(2\lambda) + 4\log 2/\pi$.

Lemma A.1.5 (Boundary series)**.** *The coefficients are real, satisfy*

$$|g_\ell^\pm| \leq M_h := \frac{Bh}{\delta} D_\lambda, \tag{A.34}$$

and give

$$\mathcal{B}_h^{\mathrm{eff}} f(x) = G^-(z_-(x)) - G^-(z_{-,0}) + G^+(z_+(x)) - G^+(z_{+,0}). \tag{A.35}$$

Proof. Expand $\coth w = 1 + 2\sum_{\ell\geq 1} e^{-2\ell w}$ and $\tanh w = 1 + 2\sum_{\ell\geq 1}(-1)^\ell e^{-2\ell w}$ for $\Re w > 0$; use oddness at the opposite endpoint. For example, writing $\Delta z_\pm^\ell = z_\pm(x)^\ell - z_{\pm,0}^\ell$, the left endpoint terms are $H_x(a+iy) = 2\sum_{\ell\geq 1}\Delta z_-^\ell e^{2i\ell\gamma y}$ and $D_x(a+iy) = \sum_{\ell\geq 1}(-1)^\ell \Delta z_-^\ell e^{2i\ell\gamma y}$; the right endpoint has a minus sign and the conjugate phase. Absolute convergence permits termwise integration, yielding (A.35). Since $\gamma d = \pi/2$,

$$\int_{-d}^{d} e^{2i\ell\gamma y}\,dy = 0 \quad (\ell \geq 1). \tag{A.36}$$

Subtract $f(a)$ or $f(b)$ in the corresponding μ integral. Cauchy's bound $|f'| \leq 2B/\delta$ gives $|\mu_\ell^\pm| \leq Bd/\delta$. Similarly, (A.23) and the integral used in the preceding lemma give $|\nu_\ell^\pm| \leq 4Bh\log 2/(\pi\delta)$. These estimates imply (A.34); conjugate symmetry gives reality. □

The cancellation makes the boundary moments $O(Bh)$ and will give an $O(Bh)$ correction readout. The quadratic exponent in the halo size instead comes from combining a correction degree proportional to R with $-\log\theta$ proportional to R.

For a series $G(z)=\sum_{\ell\ge 1} g_\ell z^\ell$, $|g_\ell|\le M_G$, use the prescribed-denominator Padé-type construction Brezinski & Redivo-Zaglia (2015):

$$D_m(z)=\prod_{i=1}^{m}(1+t_i z),\quad P_m=[D_mG]_{\le m},\quad S_m=P_m/D_m,\qquad 0<t_i\le\theta^{-1}. \tag{A.37}$$

Here $[\cdot]_{\le m}$ denotes Taylor truncation at zero and the t_i are distinct. On $|z|=\theta$, $|D_mG|\le 2^m M_G/(1-\theta)$. Cauchy's coefficient estimate, summed beyond degree m on $0\le z\le\theta^2$, and $D_m(z)\ge 1$ imply

$$\sup_{0\le z\le\theta^2}|G(z)-S_m(z)|\le\frac{M_G\theta}{(1-\theta)^2}(2\theta)^m. \tag{A.38}$$

Take $1\le m\le\lfloor(R+1)/2\rfloor$ and the outermost halo centers

$$\xi_i^-=a+(i-\tfrac12)h,\quad \xi_i^+=b-(i-\tfrac12)h,\quad t_i=e^{2\lambda(i-1/2)}. \tag{A.39}$$

Then $t_i\le\theta^{-1}$, and

$$\frac{1}{1+t_iz_-(x)}=\frac{1+\tanh(\gamma(x-\xi_i^-))}{2},\qquad \frac{1}{1+t_iz_+(x)}=\frac{1-\tanh(\gamma(x-\xi_i^+))}{2}. \tag{A.40}$$

Let $S_m^\pm$ be (A.37) for $G^\pm$ and define

$$q_{\mathrm{bd}}(x)=S_m^-(z_-(x))-S_m^-(z_{-,0})+S_m^+(z_+(x))-S_m^+(z_{+,0}),\qquad q^\star=q_h+q_{\mathrm{bd}}. \tag{A.41}$$

Proposition A.1.1 (Correction in the prescribed dictionary)**.** *The correction uses only the existing halo slots and satisfies*

$$\left\|\mathcal{B}_h^{\mathrm{eff}}f-q_{\mathrm{bd}}\right\|_\infty\le\frac{4M_h\theta}{(1-\theta)^2}(2\theta)^m. \tag{A.42}$$

With $m=\lfloor(R+1)/2\rfloor$ and $\lambda R\ge 2\log 2$,

$$\left\|\mathcal{B}_h^{\mathrm{eff}}f-q_{\mathrm{bd}}\right\|_\infty\le 16M_he^{-\lambda R^2/4}. \tag{A.43}$$

Proof. Distinct roots give $S_m=c_0+\sum_{i=1}^m c_i/(1+t_iz)$. Use (A.40) in the four anchored evaluations; constants cancel and (A.38) proves (A.42). Now $m\ge R/2$, $\theta\le e^{-\lambda R}\le 1/4$, and $(2\theta)^m\le e^{-\lambda R^2/4}$, while $4\theta/(1-\theta)^2\le 16$. □

Width-independent readout control. The partial-fraction coefficients form an inverse-Vandermonde map Gautschi (1962). For the nodes in (A.39), put

$$\zeta=e^{-2\lambda},\quad r_i=t_i^{-1}=\zeta^{i-1/2},\quad p_k=\prod_{j=1}^{k}(1-\zeta^j),\quad p_0=1. \tag{A.44}$$

The coefficients in $S_m=c_0+\sum_i c_i/(1+t_iz)$ satisfy

$$c_i=(-1)^i\frac{\zeta^{[i(i+1)-1]/2}}{p_{i-1}p_{m-i}}\sum_{\ell=1}^{m}e_{\ell-1}(r_1,\dots,\widehat{r_i},\dots,r_m)g_\ell. \tag{A.45}$$

Here e_k is the elementary symmetric polynomial and a hat omits an entry. For clarity, the specialization of the classical coefficient formula is

$$c_i=\frac{P_m(-r_i)}{\prod_{j\ne i}(1-t_j/t_i)},\qquad P_m(-r_i)=(-r_i)^m\sum_{\ell=1}^{m}g_\ell[z^{m-\ell}]\prod_{j\ne i}(1+t_jz).$$

The second equality follows by telescoping the factor $1+t_iz$ in the truncated numerator. Factoring the geometric node differences in the first equality gives (A.45). Define

$$p_\infty=\prod_{j\ge1}(1-\zeta^j),\qquad L_+=\prod_{j\ge1}(1+\zeta^{j-1/2}), \tag{A.46}$$

$$C_{\mathrm{bd}}(\lambda)=\sum_{i\ge1}\frac{\zeta^{[i(i+1)-1]/2}}{p_{i-1}p_\infty}\prod_{\substack{j\ge1\\ j\ne i}}(1+\zeta^{j-1/2}). \tag{A.47}$$

Proposition A.1.2 (Boundary coefficient bound)**.** *For fixed $\lambda > 0$, $C_{\mathrm{bd}}(\lambda) < \infty$ and*

$$\sum_{i=1}^{m} |c_i| \le C_{\mathrm{bd}}(\lambda) \sup_{\ell \ge 1} |g_\ell|, \qquad C_{\mathrm{bd}}(\lambda) \le \frac{\sqrt{\pi} e^{\lambda}}{2\sqrt{\lambda}} \exp \frac{5\pi^2}{24\lambda}. \tag{A.48}$$

Proof. In (A.45), the symmetric coefficients are nonnegative and sum to $\prod_{j \ne i}(1+r_j)$; also $p_{m-i} \ge p_\infty$. This proves the first bound. Since $\sum_j \zeta^j < \infty$, $p_\infty > 0$ and $L_+ < \infty$, and the remaining i-factor decays as $e^{-\lambda i(i+1)}$. For the explicit bound, use $p_{i-1} \ge p_\infty$ and

$$\begin{aligned} \log p_\infty^{-1} &= \sum_{k \ge 1} \frac{1}{k(e^{2\lambda k} - 1)} \le \frac{\pi^2}{12\lambda}, \\ \log L_+ &\le \frac{1}{2\lambda} \int_0^\infty \log(1 + e^{-t})\, dt = \frac{\pi^2}{24\lambda}, \\ \sum_{i \ge 1} e^{-\lambda i(i+1)} &\le \int_0^\infty e^{-\lambda t^2}\, dt = \frac{\sqrt{\pi}}{2\sqrt{\lambda}}. \end{aligned}$$

The middle inequality is the midpoint bound for the convex integrand. □

Each boundary contributes half of its partial-fraction absolute sum to the tanh readout. Thus

$$\sum_j |w_j^\star| \le \frac{2B(1+\Delta)}{\delta} + C_{\mathrm{bd}}(\lambda) \frac{Bh}{\delta} D_\lambda. \tag{A.49}$$

The anchored bias is at most $B + \sum_j |w_j^\star|$, so $h \le 1/8$ gives

$$\|c^\star\|_1 \le KB, \qquad K = 1 + \frac{4(1+\Delta)}{\delta} + \frac{C_{\mathrm{bd}}(\lambda) D_\lambda}{4\delta}. \tag{A.50}$$

The analytic tanh witness.

Theorem A.2 (Boundary-corrected analytic witness)**.** *For every admissible geometry,*

$$\|f - q^\star\|_\infty \le \frac{128\pi B}{\lambda\delta} \left[e^{-c_{\mathrm{app}} W} + \frac{e^{-\pi^2/\lambda}}{W} \right] =: E_{\mathrm{rep}}(W, \lambda), \qquad \|c^\star\|_1 \le KB. \tag{A.51}$$

In particular, $W = N + O(\sqrt{N})$ and $\gamma = \Theta(W)$ at fixed λ.

Proof. The signed decomposition is

$$f - q^\star = (\mathcal{B}_h^{\mathrm{eff}} f - q_{\mathrm{bd}}) + \mathcal{P}_h[F_x] - \mathcal{H}_h[F_x]. \tag{A.52}$$

For $A \ge \pi^2$,

$$\sum_{\ell \ge 0} \frac{1}{e^{(2\ell+1)A} - 1} \le \frac{e^{-A}}{(1 - e^{-A})(1 - e^{-2A})} \le 2e^{-A}.$$

Consequently, Lemma A.1.4 and Proposition A.1.1 give the explicit three bounds

$$\begin{aligned} E_{\mathrm{hor}} &\le \frac{32B}{\delta} e^{-\pi\delta N/4}, \\ E_{\mathrm{bd}} &\le \frac{32BD_\lambda}{\delta N} e^{-\lambda a_H^2 N/4}, \\ E_{\mathrm{repl}} &\le \frac{64\pi B}{\lambda\delta N} e^{-\pi^2/\lambda}. \end{aligned}$$

Admissibility implies $N \le W \le 2N$ and $\lambda(1 + D_\lambda) < 4\pi$. With $\kappa = \frac{1}{4}\min\{\pi\delta, \lambda a_H^2\}$, the first two terms sum to at most $(128\pi B/(\lambda\delta))e^{-\kappa W/2}$; the third is at most $(128\pi B/(\lambda\delta W))e^{-\pi^2/\lambda}$. Equation (A.50) gives the readout bound. □

A.3 Recovery and continuous error

A.3.1 Common dictionary assumptions

The recovery argument below is conditional on the following properties of a prescribed dictionary. They hold for the tanh construction just proved; the GELU/SiLU construction in Section A.4 will verify them before applying the same bounds. Assume $M \geq n$, where n is the total readout dimension, including the bias and any affine-correction units. Write the native features as ϕ_j, with $\phi_b = 1$, and set

$$\varphi_j = \phi_j/m_j, \quad \mathsf{M}_\phi = \operatorname{diag}(m_j), \quad v = \mathsf{M}_\phi c, \quad q_c = \sum_j v_j \varphi_j, \qquad m_b = 1. \tag{A.53}$$

Assume $m_j \geq 1$ and $|\varphi_j(x)| \leq 1$ on Ω. Assume also that a reference network satisfies $\|f - q^\star\|_\infty \leq E_W^{\mathrm{an}}$ and has positive deterministic coordinate envelopes

$$|v_j^\star| \leq B\alpha_j, \qquad K_W^{\mathrm{ref}} := \sum_j \alpha_j \leq K_{\mathrm{ref}}. \tag{A.54}$$

Here K_{ref} is independent of width for the fixed design.

For the continuous-error conclusions, use the core-cell Bernstein ellipses and impose the following sampling conditions. Fix $\varrho > 1$ and $p \in \mathbb{N}$, $p \geq 1$, and assume

$$\lambda(\varrho - \varrho^{-1}) \leq \pi, \qquad \frac{h}{4}(\varrho + \varrho^{-1}) < \delta, \qquad s \geq 2(p+1)^2. \tag{A.55}$$

Assume that the target and normalized features are holomorphic on a neighborhood of every closed cell ellipse, with $|f(z)| \leq B$ and $|\varphi_j(z)| \leq C_{\mathrm{ell}}$ there for a width-independent constant C_{ell}. The numerical comparison and export conditions (A.61)–(A.62), cutoff condition (A.65), and input-rounding and activation-evaluation assumptions below are additional conditions on a returned run; they do not follow from these dictionary properties.

Verification for tanh. Take $m_j = 1$ and $E_W^{\mathrm{an}} = E_{\mathrm{rep}}(W, \lambda)$ from Theorem A.2. The real-axis feature bound is immediate. Put $m = \lfloor (R+1)/2 \rfloor$ and

$$L_{i,m} = \frac{\zeta^{[i(i+1)-1]/2}}{p_{i-1}p_{m-i}} \prod_{\substack{1 \leq j \leq m \\ j \neq i}} (1 + \zeta^{j-1/2}). \tag{A.56}$$

Start each grid coordinate at $\alpha_j = h/[2(\delta - d)]$, and add $hD_\lambda L_{i,m}/(2\delta)$ at the ith corrected slot on either side. Take $\alpha_b = 1 + \sum_{j \neq b} \alpha_j$. The coefficient formulas above give (A.54); one may use $K_{\mathrm{ref}} = K$ from (A.50), because $\sum_i L_{i,m} \leq C_{\mathrm{bd}}(\lambda)$. The closed cell ellipses lie in the target tube. In preactivation coordinates their vertical semiaxis is at most $\pi/4$. For tanh,

$$|\tanh(u + iv)|^2 = \frac{\sinh^2 u + \sin^2 v}{\sinh^2 u + \cos^2 v} \leq 1 \qquad (|v| \leq \pi/4).$$

Thus the complex feature hypothesis holds with $C_{\mathrm{ell}} = 1$.

A.3.2 Reference scaling

Define

$$D = \operatorname{diag}(\sqrt{\alpha_j}), \quad L = \sqrt{\sum_j \alpha_j}, \quad S = BL, \quad A_{ij} = \varphi_j(z_i), \quad F = \frac{AD}{\sqrt{M}}, \quad y_X = \frac{(f(z_i))_{i=0}^{M-1}}{\sqrt{M}}. \tag{A.57}$$

These are exact reference quantities; rounding of their stored versions is included in the formation bounds below. Cauchy–Schwarz gives

$$a^\star = D^{-1}v^\star, \qquad \|a^\star\|_2 \leq S, \quad \|Fa^\star - y_X\|_2 \leq E_W^{\mathrm{an}}, \quad \|Da\|_1 \leq L\|a\|_2. \tag{A.58}$$

Indeed, $\sum_j |v_j^\star|^2/\alpha_j \leq B^2 \sum_j \alpha_j$; the residual bound follows from uniform approximation.

A.3.3 NUMERICAL CONDITIONS AND COMPARISON PROBLEM

Use round-to-nearest arithmetic with unit roundoff u, gradual underflow, and no overflow or invalid operations in a returned run. Let μ be a valid absolute rounding allowance (the smallest positive subnormal is sufficient), so $|\,\mathrm{RN}(t) - t| \le u|t| + \mu$. All error bounds are either standard a priori bounds or validated upper bounds with outward rounding. Failed range or denominator checks do not produce a certified return. Standard product, summation, and matrix-operation estimates are taken from (Higham, 2002, Chs. 2–4); SVD factorization and orthogonality bounds are supplied by (Anderson et al., 1999, Ch. 4).

For later use, a length-q pairwise dot product with rounded products obeys

$$|\,\mathrm{pdot}(v,w) - v^T w| \le \Gamma_q^{\mathrm{pw}} \sum_i |v_i w_i| + \beta_q^{\mathrm{pw}}, \quad \Gamma_q^{\mathrm{pw}} = \frac{k_q u}{1 - k_q u}, \quad \beta_q^{\mathrm{pw}} = \frac{(2q-1)\mu}{1 - k_q u}, \quad k_q = 1 + \lceil \log_2 q \rceil. \tag{A.59}$$

Assume $k_q u < 1$. The relative bound is the usual pairwise bound; the additive term counts at most $2q - 1$ local absolute errors and their amplification along the tree. In particular $\Gamma_q^{\mathrm{pw}} = O(u \log q)$.

The algorithm assembles $\widetilde{F}, \widetilde{y}$, computes a thin SVD $(\widehat{U}, \widehat{\Sigma}, \widehat{V})$ with nonnegative diagonal $\widehat{\Sigma}$, forms $p_i = \mathrm{pdot}(\widehat{U}_{:i}, \widetilde{y})$, and retains exactly the singular values $\widehat{\sigma}_i > \tau$. It sets

$$\widehat{z}_i = \begin{cases} \mathrm{RN}(p_i/\widehat{\sigma}_i), & \widehat{\sigma}_i > \tau, \\ 0, & \widehat{\sigma}_i \le \tau, \end{cases} \tag{A.60}$$

and exports the native readout by forming $D\widehat{V}\widehat{z}$ and dividing coordinate j by m_j, all in the working arithmetic. The bias is included in this same solve.

We require bounds for the exact comparison quantities

$$F_0 = U_0 \widehat{\Sigma} V_0^T, \quad U_0^T U_0 = V_0^T V_0 = I, \qquad a_\tau = F_{0,\tau}^{\dagger} y_0 = V_0 \widehat{z}, \tag{A.61}$$

with

$$\|F - F_0\|_2 \le a_W, \quad \|y_X - y_0\|_2 \le b_W, \quad \|\mathsf{M}_\phi \widehat{c} - D a_\tau\|_1 \le e_{c,W}. \tag{A.62}$$

The matrix and label formation bounds include normalization, evaluation at stored sample points, and any hidden-geometry storage error. Thus the analytic basis remains the prescribed basis, rather than an implicitly perturbed one.

Here is a sufficient realization of (A.61)–(A.62) that also makes the dependence on the cutoff explicit. Suppose

$$\left\|F - \widetilde{F}\right\|_2 \le e_F, \qquad \|y_X - \widetilde{y}\|_2 \le e_y, \qquad \left\|\widetilde{F} - \widehat{U}\widehat{\Sigma}\widehat{V}^T\right\|_2 \le r_F,$$
$$\left\|\widehat{U}^T\widehat{U} - I\right\|_2 \le o_U < 1, \qquad \left\|\widehat{V}^T\widehat{V} - I\right\|_2 \le o_V < 1.$$

Put $\delta_U = 1 - \sqrt{1 - o_U}$, $\delta_V = 1 - \sqrt{1 - o_V}$, $U_+ = \sqrt{1 + o_U}$, $V_+ = \sqrt{1 + o_V}$, and $Y = \|\widetilde{y}\|_2$. Standard bounds or enclosed residual products supply these numbers; spectral norms may be bounded by Frobenius norms. Then one may take

$$\begin{aligned} a_W &= e_F + r_F + \left\|\widehat{\Sigma}\right\|_2 (\delta_U V_+ + \delta_V), \\ e_p &= \sqrt{n}\,[\Gamma_M^{\mathrm{pw}} U_+ Y + \beta_M^{\mathrm{pw}}], \\ b_W &= e_y + \delta_U Y + e_p + u\,\|p\|_2 + \mu\,\|\widehat{\sigma}\|_2 . \end{aligned} \tag{A.63}$$

To verify the interface, take the polar factors $U_0 = \widehat{U}(\widehat{U}^T\widehat{U})^{-1/2}$ and similarly V_0. Their singular values give $\left\|\widehat{U} - U_0\right\|_2 \le \delta_U$ and $\left\|\widehat{V} - V_0\right\|_2 \le \delta_V$, proving the bound for a_W. If $J_\tau = \{i : \widehat{\sigma}_i > \tau\}$, set

$$y_0 = \widetilde{y} + U_{0,J_\tau}(\widehat{\Sigma}_{J_\tau}\widehat{z}_{J_\tau} - U_{0,J_\tau}^T \widetilde{y}).$$

This makes $F_{0,\tau}^{\dagger} y_0 = V_0 \widehat{z}$ exactly. Projection rounding contributes e_p, and division gives $|\widehat{\sigma}_i \widehat{z}_i - p_i| \le u|p_i| + \mu\widehat{\sigma}_i$ on retained coordinates. These bounds prove the formula for b_W. In particular,

a_W and b_W are available before choosing τ; the bound for division does not contain $1/\tau$. The polar factors and y_0 are analysis objects, not computed intermediates.

Finally let $e_{\exp}$ be the standard forward-error bound for the specified native export, in the form

$$\left\| \mathsf{M}_\phi \widehat{c} - D\widehat{V}\widehat{z} \right\|_1 \le e_{\exp}.$$

It includes storage of D, row dot products, products, and division by m_j. It can also be obtained by enclosing each displayed coordinate difference using (A.59) and the scalar rounding model. Once the bound $\|\widehat{z}\|_2 \le 2S$ below is established, the choice

$$e_{c,W} = e_{\exp} + 2LS\delta_V \tag{A.64}$$

proves the last inequality of (A.62). This construction of the numerical interface uses only standard error bounds and returned finite arrays; it does not assume recovery of the analytic coefficients.

A.3.4 Cutoff and recovery bound

Let $\overline{E}_W \ge E_W^{\mathrm{an}}$ be the approximation bound available to the algorithm, and retain its formation allowance $\epsilon_{E,W} := \overline{E}_W - E_W^{\mathrm{an}} \ge 0$. Choose a positive threshold with

$$\eta_W = \overline{E}_W + a_W S + b_W, \qquad \tau S \ge \eta_W, \qquad 0 \le \tau S - \eta_W \le \rho_W. \tag{A.65}$$

An outward upper bound for η_W/S provides such a threshold; ρ_W includes this excess and its rounding. No condition on singular-value gaps is needed. The truncated-SVD coefficient comparison in (Adcock & Huybrechs, 2020, Theorem 3.8, first inequality) gives

$$\|a_\tau\|_2 \le S + \eta_W/\tau \le 2S, \qquad \|F_0 a_\tau - y_0\|_2 \le \eta_W + \tau S. \tag{A.66}$$

For the second inequality, the discarded singular values are at most τ; apply this to the comparison vector $a^\star$. This is the usual thresholded least-squares estimate Hansen (1987).

Proposition A.2.1 (Recovered coefficient and sample bounds)**.** *Define*

$$T = 2LS + e_{c,W}, \qquad R_X = 2\overline{E}_W + 4a_W S + 3b_W + \rho_W + e_{c,W}. \tag{A.67}$$

The returned native readout satisfies

$$\|\widehat{c}\|_1 \le \|\mathsf{M}_\phi \widehat{c}\|_1 \le T, \qquad \|f - q_{\widehat{c}}\|_{2,X} \le R_X. \tag{A.68}$$

Proof. The reference residual on (F_0, y_0) is at most η_W, so (A.66) applies. Since $a_\tau = V_0\widehat{z}$, $\|\widehat{z}\|_2 = \|a_\tau\|_2 \le 2S$, which justifies (A.64). Returning to F, y_X gives

$$\|Fa_\tau - y_X\|_2 \le \eta_W + \tau S + 2a_W S + b_W \le 2\overline{E}_W + 4a_W S + 3b_W + \rho_W.$$

Now use (A.58), coefficient export, and $|\varphi_j| \le 1$. Finally, $m_j \ge 1$ bounds the native norm by the normalized norm. The reference part of T is $2BL^2 \le 2BK_{\mathrm{ref}}$. □

At fixed observation density, the pairwise projection contribution to e_p alone is $O(u\sqrt{W}\log W)$ times $U_+ Y$.

A.3.5 From sampled error to continuous error

Let $q = q_{\widehat{c}}$ be the exact function of the returned readout in the prescribed basis, and put $g = f - q$. For each core cell $I = [x_k, x_{k+1}]$, define

$$\|v\|_{2,I}^2 = \frac{1}{h}\int_I |v|^2, \qquad \|v\|_{\mathrm{trap},I}^2 = \frac{1}{s}\left[\frac{|v(x_k)|^2 + |v(x_{k+1})|^2}{2} + \sum_{j=1}^{s-1} |v(x_k + jh/s)|^2\right].$$

Under the common dictionary assumptions, g is holomorphic on the cell ellipses and the recovered normalized coefficient bound gives

$$|g(z)| \le \mathcal{M} := B + C_{\mathrm{ell}} T \quad \text{on every cell ellipse.} \tag{A.69}$$

The analytic Chebyshev estimate (Trefethen, 2019, Ch. 8) gives a degree-p polynomial P_I such that

$$\|g - P_I\|_{\infty,I} \le e_p^{\mathrm{an}} := \frac{2\mathcal{M}}{\varrho - 1}\varrho^{-p}. \tag{A.70}$$

The local polynomial inverse estimates used below are standard; see (Canuto et al., 2006, Secs. 5.4.1 and 5.4.4). We record their constants and the specialization to the present equispaced grid.

Lemma A.2.1 (Local polynomial sampling)**.** *For every polynomial P of degree at most p on a cell of length h,*

$$\|P'\|_{2,I} \le \frac{(p+1)\sqrt{p(p+2)}}{h} \|P\|_{2,I}\,, \tag{A.71}$$

$$\|P\|_{\infty,I} \le (p+1)\,\|P\|_{2,I}\,, \qquad \|P'\|_{\infty,I} \le \frac{2p^2(p+1)}{h}\,\|P\|_{2,I}\,. \tag{A.72}$$

Under (A.55),

$$\|P\|_{2,I} \le \sqrt{2}\,\|P\|_{\mathrm{trap},I}\,, \qquad \|P\|_{\infty,I} \le 2 \max_{x\in X\cap I} |P(x)|. \tag{A.73}$$

Proof. For completeness, the constants in the cited inverse estimates can be obtained from the normalized Legendre basis $\sqrt{2j+1}\,P_j$ on $[-1,1]$: $|P_j| \le 1$, $\int_{-1}^{1}(P_j')^2 = j(j+1)$, and $\sum_{j=1}^{p}(2j+1)j(j+1)/2 = p(p+1)^2(p+2)/4$. Cauchy–Schwarz and rescaling give (A.71) and the first inequality in (A.72); Markov's inequality gives the second. For the equispaced grid, the elementary trapezoidal error for $|P|^2$ gives

$$\left|\|P\|_{\mathrm{trap},I}^2 - \|P\|_{2,I}^2\right| \le \frac{h}{s}\,\|P\|_{2,I}\,\|P'\|_{2,I} \le \frac{(p+1)\sqrt{p(p+2)}}{s}\,\|P\|_{2,I}^2 \le \tfrac{1}{2}\,\|P\|_{2,I}^2\,.$$

This proves the first part of (A.73). For the uniform norm, choose the nearest grid point to a maximizer, at distance at most $h/(2s)$, and use Markov's inequality to get $\|P\|_\infty \le \max_X |P| + (p^2/s)\,\|P\|_\infty$. □

Lemma A.2.2 (Continuous error transfer)**.** *With $c_M = \sqrt{2M/(M-1)} \le 2$,*

$$\|g\|_{2,\Omega} \le c_M\,\|g\|_{2,X} + (1+\sqrt{2})e_p^{\mathrm{an}}, \tag{A.74}$$

$$\|g\|_\infty \le 2\max_{x\in X}|g(x)| + 3e_p^{\mathrm{an}}. \tag{A.75}$$

Proof. On each cell, (A.70) and (A.73) give $\|g\|_{2,I} \le \sqrt{2}\,\|g\|_{\mathrm{trap},I} + (1+\sqrt{2})e_p^{\mathrm{an}}$. Combine cells by Minkowski. The global trapezoidal norm squared is

$$\frac{1}{M-1}\left[\frac{|g(z_0)|^2 + |g(z_{M-1})|^2}{2} + \sum_{i=1}^{M-2}|g(z_i)|^2\right] \le \frac{M}{M-1}\,\|g\|_{2,X}^2\,.$$

For the uniform inequality, apply the second part of (A.73) to P_I, and add the three approximation-error terms. □

Input rounding and evaluation. Let $Q(x)$ be round-to-nearest in the working format, clamped to Ω, and let $\xi \ge \sup_{x\in\Omega}|Q(x)-x|$. For the RMS rounding estimate assume that the core-cell endpoints are exactly representable. Monotonicity of round-to-nearest then gives $Q(I) \subseteq I$ for each core cell. For example, $N = 2^k$ satisfies this condition when the required grid significands, exponents, and integer indices are exactly representable. Set

$$\chi_W = \frac{2p^2(p+1)\xi}{h}. \tag{A.76}$$

The bound on P_I' in (A.72) implies $\|P_I \circ Q\|_{2,I} \le (1+\chi_W)\,\|P_I\|_{2,I}$. Repeating the proof of Lemma A.2.2 therefore gives

$$\|g\circ Q\|_{2,\Omega} \le (1+\chi_W)c_M R_X + [1+\sqrt{2}(1+\chi_W)]e_p^{\mathrm{an}}. \tag{A.77}$$

Also $\|f - f\circ Q\|_{2,\Omega} \le B\xi/\delta$ by Cauchy's estimate. We collect the rounding terms separately:

$$E_{\mathrm{in,W}} = \frac{B\xi}{\delta} + \chi_W(c_M R_X + \sqrt{2}e_p^{\mathrm{an}}). \tag{A.78}$$

Assume the activation evaluator, including stored affine arguments, satisfies

$$|\widehat{\phi}_j(t) - \phi_j(t)| \le m_j\epsilon_{\phi,W} \quad (t \in \Omega \text{ representable}).$$

This is a uniform error contract for the specified tanh, exact-GELU, or SiLU routine, including any geometry-storage perturbation; it is not an assumption of an exact transcendental evaluation. Use rounded products and pairwise summation for the native readout. Since $\sum_j |\widehat{c}_j \phi_j(t)| \le T$ and $\sum_j |\widehat{c}_j \widehat{\phi}_j(t)| \le (1+\epsilon_{\phi,W})T$, (A.59) gives

$$E_{\mathrm{ev,W}} := [\epsilon_{\phi,W} + \Gamma_n^{\mathrm{pw}}(1+\epsilon_{\phi,W})]T + \beta_n^{\mathrm{pw}}, \qquad |\widehat{q}_W(x) - q(Q(x))| \le E_{\mathrm{ev,W}}. \tag{A.79}$$

The mixed term $\Gamma_n^{\mathrm{pw}}\epsilon_{\phi,W}$ is retained.

Theorem A.3 (Computed continuous-error bounds)**.** *Under the preceding analytic and numerical conditions,*

$$\|\widehat{c}_W\|_1 \le T, \qquad \|f - \widehat{q}_W\|_{2,\Omega} \le c_M R_X + (1+\sqrt{2})e_p^{\mathrm{an}} + E_{\mathrm{in,W}} + E_{\mathrm{ev,W}}. \tag{A.80}$$

If $r_{\infty,X}$ is a validated upper bound on $\max_{x\in X}|f(x)-q(x)|$, then also

$$\|f - \widehat{q}_W\|_\infty \le 2r_{\infty,X} + 3e_p^{\mathrm{an}} + B\xi/\delta + E_{\mathrm{ev,W}}. \tag{A.81}$$

The uniform statement requires only $Q(\Omega) \subseteq \Omega$, not cellwise preservation. It is a maximum-residual certificate, not a replacement of the sampled RMS by the same numerical value.

Proof. The coefficient bound is Proposition A.2.1. Decompose $f - \widehat{q} = (f - f\circ Q) + (g\circ Q) + (q\circ Q - \widehat{q})$, and apply (A.77) and (A.79) for RMS. For the uniform norm, use (A.75) at $Q(x)$. □

A.3.6 Tanh theorem and width dependence

Put

$$A_{\mathrm{rep}} = \frac{128\pi}{\lambda\delta}, \qquad \Delta_{\mathrm{rec,W}} = 2\epsilon_{E,W} + 4a_W S + 3b_W + \rho_W + e_{c,W}. \tag{A.82}$$

For tanh, $R_X = 2E_{\mathrm{rep}}(W,\lambda) + \Delta_{\mathrm{rec,W}}$. Using $c_M \le 2$ in the analytic term of (A.80) gives

$$\|f - \widehat{q}_W\|_{2,\Omega} \le 4A_{\mathrm{rep}}B\left(e^{-c_{\mathrm{app}}W} + \frac{e^{-\pi^2/\lambda}}{W}\right) + \epsilon_{\mathrm{rec,W}}, \tag{A.83}$$

where

$$\epsilon_{\mathrm{rec,W}} = c_M\Delta_{\mathrm{rec,W}} + (1+\sqrt{2})e_p^{\mathrm{an}} + E_{\mathrm{in,W}} + E_{\mathrm{ev,W}}, \qquad \epsilon_{\mathrm{coeff,W}} = e_{c,W}. \tag{A.84}$$

These quantities are nonnegative explicit bounds; $E_{\mathrm{in,W}}$ contains all factors involving ξ/h. The reference coefficient bound is $2LS = 2BK_W^{\mathrm{ref}} \le 2BK_{\mathrm{ref}}$, which completes the proof of Theorem A.1.

Corollary A.3.1 (Logarithmic sufficient width above a fixed allowance)**.** *Fix the target class, bandwidth, activation, and design. Suppose the preceding bounds have the form*

$$\|f - \widehat{q}_W\|_{2,\Omega} \le Ae^{-cW} + \tau_W$$

with $A, c > 0$ independent of width, and $\tau_W \le \tau$ throughout a specified certified range $\mathcal{W}$ of admissible widths. For any $\epsilon > \tau$, every $W \in \mathcal{W}$ satisfying

$$W \ge \frac{1}{c}\log\frac{A}{\epsilon - \tau} \tag{A.85}$$

has error at most ϵ. For tanh one may take $A = 4A_{\mathrm{rep}}B$, $c = c_{\mathrm{app}}$, and $\tau_W = 4A_{\mathrm{rep}}BW^{-1}e^{-\pi^2/\lambda} + \epsilon_{\mathrm{rec,W}}$.

Proof. The error is bounded by $Ae^{-cW} + \tau$; solve $Ae^{-cW} \le \epsilon - \tau$. Existence of an admissible width meeting this inequality inside the certified range remains part of the premise. □

For $\epsilon \ge 2\tau$ the sufficient width is $c^{-1}\log(2A/\epsilon)$, subject to the same range restriction. This is the fixed-design statement underlying Corollary 3.1.1. If λ is instead varied with ϵ, the dependence of A, c, admissibility, and the allowances on λ must be retained; this corollary alone makes no joint $\epsilon \to 0$ assertion with varying design.

A.4 Extensions to exact GELU and SiLU

For these activations we approximate f'' by a polynomial, apply a finite differential inverse of the kernel convolution, and discretize by Lemma A.1.3. We then verify the dictionary assumptions of Section A.3 and apply its recovery and continuous-error bounds.

A.4.1 Domains and polynomial approximation

Choose $1 < L_* < 1 + \Delta$ and let $J_* = [-L_*, L_*]$. Denote by $\mathcal{E}_\rho(J_*)$ the closed filled Bernstein ellipse with semiaxes $L_*(\rho + \rho^{-1})/2$ and $L_*(\rho - \rho^{-1})/2$. Fix $1 < \rho_1 < \rho_2$, $\delta_* > 0$, and $s_* > 0$ such that

$$\mathcal{E}_{\rho_2}(J_*) + \{z : |z| \le \delta_*\} \subset U_\delta(I_\Delta), \qquad \mathcal{D} := J_* + i[-s_*, s_*] \Subset \operatorname{int} \mathcal{E}_{\rho_1}(J_*). \tag{A.86}$$

Such choices exist by taking the ellipses and s_*, δ_* sufficiently small. Set

$$d_* := \operatorname{dist}(\mathcal{D}, \partial\mathcal{E}_{\rho_1}(J_*)) > 0, \quad d_{\mathrm{ext}} = L_* - 1, \quad C_* := 2/d_{\mathrm{ext}}. \tag{A.87}$$

Cauchy's estimate bounds f'' on $\mathcal{E}_{\rho_2}(J_*)$ by $2B/\delta_*^2$. The analytic Chebyshev estimates of (Trefethen, 2019, Ch. 8), applied to its degree-ν truncation, give a real polynomial p_ν with

$$\|f'' - p_\nu\|_{L^\infty(J_*)} \le A_{\mathrm{pol}} B \rho_2^{-\nu}, \qquad \sup_{\mathcal{E}_{\rho_1}(J_*)} |p_\nu| \le A_{\mathrm{pol}} B, \tag{A.88}$$

where, for example, one may take

$$A_{\mathrm{pol}} = \frac{2}{\delta_*^2} \max\left\{ \frac{2}{\rho_2 - 1},\ 1 + \frac{2\rho_1}{\rho_2 - \rho_1} \right\}. \tag{A.89}$$

The second inequality follows by summing the Chebyshev coefficient bound on the smaller ellipse. We reserve n for the number of readout coordinates and use $\nu = \lfloor \eta N \rfloor$ for the polynomial degree.

A.4.2 Finite polynomial deconvolution

For a real even kernel K with an exponential absolute moment and integral one, write

$$M_K(z) = \int_{\mathbb{R}} K(t) e^{zt}\, dt, \qquad M_K(z)^{-1} = \sum_{\ell \ge 0} b_{2\ell} z^{2\ell}. \tag{A.90}$$

Suppose the inverse is holomorphic on $|z| \le r_*$ and bounded there by C_{inv}. Define

$$a_{\gamma,\nu} = \sum_{\ell=0}^{\lfloor \nu/2 \rfloor} b_{2\ell} \gamma^{-2\ell} p_\nu^{(2\ell)}, \qquad A_{\mathrm{den}} := \tfrac{4}{3} C_{\mathrm{inv}} A_{\mathrm{pol}}. \tag{A.91}$$

Lemma A.3.1 (Polynomial deconvolution)**.** *The density is a real polynomial. With* $K_\gamma(t) = \gamma K(\gamma t)$*,*

$$\int_{\mathbb{R}} a_{\gamma,\nu}(t) K_\gamma(x - t)\, dt = p_\nu(x). \tag{A.92}$$

If $\nu/(r_* \gamma d_*) \le 1/2$*, then*

$$\sup_{z \in \mathcal{D}} |a_{\gamma,\nu}(z)| \le A_{\mathrm{den}} B. \tag{A.93}$$

For instance, $0 < \eta \le r_* \lambda d_*/4$ *ensures this degree condition.*

Proof. Taylor expansion of any polynomial a terminates, so evenness of K gives

$$\int_{\mathbb{R}} K(u) a(x - u/\gamma)\, du = M_K(D/\gamma) a(x), \qquad D = d/dx.$$

All integrals exist by the moment hypothesis. Since D is nilpotent on polynomials of degree at most ν, the finite substitution of the inverse series gives (A.92) exactly. Cauchy's estimates give $|b_{2\ell}| \le C_{\mathrm{inv}} r_*^{-2\ell}$ and $|p_\nu^{(2\ell)}(z)| \le A_{\mathrm{pol}} B (2\ell)!/d_*^{2\ell}$ for $z \in \mathcal{D}$. Using $(2\ell)! \le \nu^{2\ell}$ and summing a geometric series of ratio at most $1/4$ proves (A.93). □

We will also need an exterior bound, since the convolution in (A.92) is over the whole real line. The real Chebyshev extremal inequality gives, for $|t| \geq L_*$,

$$|a_{\gamma,\nu}(t)| \leq A_{\mathrm{den}} B \exp\left(\nu \operatorname{arcosh} \frac{|t|}{L_*}\right) \leq A_{\mathrm{den}} B e^{\nu C_*(|t|-1)}. \tag{A.94}$$

The extremal inequality follows from alternation at the $\nu + 1$ Chebyshev extrema: a larger value outside J_* would force the difference from a suitable multiple of $T_\nu(t/L_*)$ to have more than ν zeros. The last inequality uses $\operatorname{arcosh} s \leq \log(2s) \leq 2s$ and $|t|/L_* \leq (|t|-1)/(L_*-1)$.

A.4.3 Kernel identities and admissibility

Let

$$\varphi(z) = \frac{e^{-z^2/2}}{\sqrt{2\pi}}, \qquad \Phi(z) = \tfrac{1}{2} + \int_0^z \varphi(t)\,dt, \qquad \varsigma(z) = (1+e^{-z})^{-1},$$

$$\psi_G(z) = z\Phi(z), \qquad \psi_S(z) = z\varsigma(z).$$

These are exact GELU and SiLU. Differentiation gives

$$K_G(z) = \psi_G''(z) = (2-z^2)\frac{e^{-z^2/2}}{\sqrt{2\pi}}, \qquad M_G(z) = (1-z^2)e^{z^2/2}, \tag{A.95}$$

$$K_S(z) = \psi_S''(z) = 2g(z) + zg'(z), \quad g = \varsigma' = \tfrac{1}{4}\operatorname{sech}^2(z/2), \qquad M_S(z) = \frac{(\pi z)^2\cos(\pi z)}{\sin^2(\pi z)}. \tag{A.96}$$

At zero, both moment functions have removable value one. For GELU the moment identity follows by differentiating $\int \varphi(t)e^{zt}\,dt = e^{z^2/2}$ twice. For SiLU, the substitution $v = e^t$ gives

$$G(z) := \int_{\mathbb{R}} g(t)e^{zt}\,dt = \int_0^\infty \frac{v^z}{(1+v)^2}\,dv = \Gamma(1+z)\Gamma(1-z) = \frac{\pi z}{\sin \pi z} \quad (|\Re z| < 1).$$

Integration by parts in $\int tg'(t)e^{zt}\,dt$ yields $M_S = G - zG'$, proving (A.96). Thus

$$M_G^{-1}(z) = \frac{e^{-z^2/2}}{1-z^2}, \qquad M_S^{-1}(z) = \frac{\sin^2(\pi z)}{(\pi z)^2\cos(\pi z)}. \tag{A.97}$$

The inverse radii are 1 and $1/2$, respectively. For $0 < r_* < 1$ in the GELU case and $0 < r_* < 1/2$ in the SiLU case, valid bounds are

$$C_{\mathrm{inv},G} = \frac{e^{r_*^2/2}}{1-r_*^2}, \qquad C_{\mathrm{inv},S} = \frac{[\sinh(\pi r_*)/(\pi r_*)]^2}{\cos(\pi r_*)}. \tag{A.98}$$

These follow from the power series of sine and $|\cos(\pi z)|^2 = \cos^2(\pi\Re z) + \sinh^2(\pi\Im z)$.

Choose $a_G > 0$, $0 < a_S < d_{\mathrm{ext}}/2$, and $0 < \vartheta < \pi$, and set

$$R_G = \lceil a_G\sqrt{N}\rceil, \quad R_S = \lceil a_S N\rceil, \quad r_R = (R+\tfrac{1}{2})h, \quad \sigma_G = \frac{2\pi h}{\lambda^2}, \quad \sigma_S = \frac{\vartheta}{\gamma}. \tag{A.99}$$

For each activation choose $\eta > 0$ as in Lemma A.3.1, with $\eta \leq 1$. In the SiLU case also require $\eta \leq \lambda/(8C_*)$. In this subsection take N large enough that $\nu \geq 1$, $\gamma \geq 1$, and

$$r_R \leq d_{\mathrm{ext}}, \qquad \sigma_\psi \leq s_*, \qquad \begin{cases} \nu C_* \leq \gamma^2 d_{\mathrm{ext}}/4, & \psi = \psi_G, \\ \nu C_* \leq \gamma/4, & \psi = \psi_S. \end{cases} \tag{A.100}$$

All these conditions hold for sufficiently large N at fixed design. In particular, the midpoint rectangle is contained in $\mathcal{D}$. The SiLU contour stays below its nearest poles at preactivation height π.

A.4.4 HALO AND MIDPOINT ESTIMATES

Let $J_R = [a_R, b_R] = [-1 - r_R, 1 + r_R]$ and write

$$(\mathcal{C}_R p_\nu)(x) = \int_{J_R} a_{\gamma,\nu}(t) K_\gamma(x - t)\, dt,$$
$$E^\psi_{\mathrm{halo}}(x) = \int_{\mathbb{R}\setminus J_R} a_{\gamma,\nu}(t) K_\gamma(x - t)\, dt, \quad F^\psi_x(z) = a_{\gamma,\nu}(z) K_\gamma(x - z),$$
$$G^\psi_N(x) = h \sum_{j=-R}^{N+R} a_{\gamma,\nu}(x_j) K_\gamma(x - x_j). \tag{A.101}$$

There are no kernel poles in the chosen rectangles. Hence Lemma A.1.3 and (A.92) give

$$f'' - G^\psi_N = (f'' - p_\nu) + E^\psi_{\mathrm{halo}} - \mathcal{H}_h[F^\psi_x] - \mathcal{V}_h[F^\psi_x]. \tag{A.102}$$

Put $Z = \gamma r_R = \lambda(R + 1/2)$. The following estimates include explicit prefactors for the terms that will enter the approximation bound.

Proposition A.3.1 (GELU kernel quadrature)**.** *Define*

$$Q_G(\lambda) = \frac{2(3 + 4\pi^2/\lambda^2) e^{-2\pi^2/\lambda^2}}{1 - e^{-4\pi^2/\lambda^2}},$$
$$V_G(\lambda) = \frac{4}{\sqrt{2\pi}} \left(\frac{4\lambda}{\pi} + \frac{2\lambda^3}{\pi^3} \right). \tag{A.103}$$

Under (A.100), uniformly on Ω*,*

$$|E^G_{\mathrm{halo}}| \le 12 A_{\mathrm{den}} B e^{-Z^2/8}, \quad |\mathcal{H}_h[F^G_x]| \le A_{\mathrm{den}} B Q_G(\lambda), \quad |\mathcal{V}_h[F^G_x]| \le A_{\mathrm{den}} B V_G(\lambda) e^{-Z^2/4}. \tag{A.104}$$

Proof. For real u, $|K_G(u)| \le (2 + u^2) e^{-u^2/2}/\sqrt{2\pi}$. On $J_* \setminus J_R$, use the bounded density. Outside J_*, let $w = |t - x| \ge d_{\mathrm{ext}}$. Equations (A.94) and (A.100) imply $|a_{\gamma,\nu}(t)| \le A_{\mathrm{den}} B e^{\gamma^2 w^2/4}$. Thus on both omitted tails, with $u = \gamma w$, the absolute integrand is at most $A_{\mathrm{den}} B \gamma (2 + u^2) e^{-u^2/4}/\sqrt{2\pi}$. Integration gives

$$|E^G_{\mathrm{halo}}| \le \frac{2 A_{\mathrm{den}} B}{\sqrt{2\pi}} \int_Z^\infty (2 + u^2) e^{-u^2/4}\, du \le 12 A_{\mathrm{den}} B e^{-Z^2/8},$$

using $\int_0^\infty (2 + u^2) e^{-u^2/8}\, du = 6\sqrt{2\pi}$.

For the horizontal sides, for any real v,

$$\int_{\mathbb{R}} |K_G(u + iv)|\, du \le (3 + v^2) e^{v^2/2}.$$

Choose $v_G = \gamma \sigma_G = 2\pi/\lambda$ and use (A.21). The two horizontal sides contribute at most $2 A_{\mathrm{den}} B (3 + v_G^2) e^{v_G^2/2}/(e^{2\pi v_G/\lambda} - 1)$, which is $A_{\mathrm{den}} B Q_G$.

On a vertical side, the real kernel argument obeys $|u| \ge Z$. The elementary inequality $(2 + u^2 + v^2) e^{-u^2/2} \le (4 + v^2) e^{-Z^2/4}$ gives

$$|\mathcal{V}_h[F^G_x]| \le \frac{4 A_{\mathrm{den}} B}{\sqrt{2\pi}} e^{-Z^2/4} \int_0^{v_G} \frac{(4 + v^2) e^{v^2/2}}{1 + e^{2\pi v/\lambda}}\, dv.$$

For $0 \le v \le v_G$, the exponential ratio is at most $e^{-\pi v/\lambda}$. Extend the integral to infinity to obtain V_G. □

Proposition A.3.2 (SiLU kernel quadrature)**.** *Let* $c_\vartheta = \sec(\vartheta/2)$ *and define*

$$A_\vartheta = 2c_\vartheta^2 + (1 + \vartheta) c_\vartheta^3, \quad Q_S(\lambda) = \frac{8 A_\vartheta}{e^{2\pi\vartheta/\lambda} - 1}, \quad V_S(\lambda) = \frac{4\lambda A_\vartheta \log 2}{\pi}. \tag{A.105}$$

Under (A.100), uniformly on Ω*,*

$$|E^S_{\mathrm{halo}}| \le 16 A_{\mathrm{den}} B e^{-Z/2}, \quad |\mathcal{H}_h[F^S_x]| \le A_{\mathrm{den}} B Q_S(\lambda), \quad |\mathcal{V}_h[F^S_x]| \le A_{\mathrm{den}} B V_S(\lambda) e^{-Z/2}. \tag{A.106}$$

Proof. On the real axis $|g(u)| \le e^{-|u|}$ and $g' = g(1-2\varsigma)$, so $|K_S(u)| \le (2+|u|)e^{-|u|}$. Outside J_*, (A.94) and $\nu C_* \le \gamma/4$ bound the density by $A_{\mathrm{den}} B e^{\gamma|t-x|/4}$. Within J_* use its uniform bound. Consequently

$$|E^S_{\mathrm{halo}}| \le 2A_{\mathrm{den}} B \int_Z^\infty (2+u)e^{-3u/4}\,du \le 16 A_{\mathrm{den}} B e^{-Z/2},$$

since $(2+u)e^{-u/4} \le 4$ for $u \ge 0$. For $z = u+iv$, $|v| \le \vartheta$, writing $t = e^{-u}$ gives

$$|1+te^{-iv}| \ge (1+t)\cos(\vartheta/2), \quad |g(z)| \le c_\vartheta^2 e^{-|u|}, \quad |1-2\varsigma(z)| \le c_\vartheta.$$

Therefore $|K_S(u+iv)| \le A_\vartheta(1+|u|)e^{-|u|}$, whose real-line integral is at most $4A_\vartheta$. The two horizontal sides and (A.21) give Q_S. On each vertical side $|u| \ge Z$, and $(1+|u|)e^{-|u|} \le 2e^{-|u|/2}$. Four endpoint terms and $\int_0^\infty (1+e^{2\pi y/h})^{-1} dy = h\log 2/(2\pi)$ give V_S. □

A.4.5 Activation realization and coefficient envelopes

Define

$$q_0^\psi(x) = \frac{h}{\gamma} \sum_{j=-R}^{N+R} a_{\gamma,\nu}(x_j)\psi(\gamma(x-x_j)). \tag{A.107}$$

Twice differentiating gives $(q_0^\psi)'' = G_N^\psi$. Choose

$$\beta = f'(-1) - (q_0^\psi)'(-1), \quad \alpha = f(-1) - q_0^\psi(-1) + \beta, \quad q_\psi^\star = q_0^\psi + \alpha + \beta x. \tag{A.108}$$

Both activations obey $\psi(z) - \psi(-z) = z$ on their domains, by $\Phi(-z) = 1-\Phi(z)$ or $\varsigma(-z) = 1-\varsigma(z)$. Thus

$$\beta x = \frac{\beta}{\gamma}[\psi(\gamma x) - \psi(-\gamma x)] \tag{A.109}$$

uses two additional hidden units and α is the output bias. The two anchor conditions imply

$$f(x) - q_\psi^\star(x) = \int_{-1}^x (x-t)[f''(t) - G_N^\psi(t)]\,dt. \tag{A.110}$$

Integration gives a factor at most two in L^∞; Young's inequality with $s\mathbf{1}_{[0,2]}(s)$ gives the same factor in normalized L^2.

Theorem A.4 (Exact-GELU and SiLU witnesses)**.** *With the domains, parameters, and admissibility conditions above, the prescribed network has $W_\psi = N + 2R_\psi + 3$ hidden-unit slots and satisfies*

$$\|f - q_G^\star\|_\infty \le 2B\left[A_{\mathrm{pol}}\rho_2^{-\nu} + A_{\mathrm{den},G}\{(12+V_G)e^{-Z_G^2/8} + Q_G\}\right] =: E_W^G, \tag{A.111}$$

$$\|f - q_S^\star\|_\infty \le 2B\left[A_{\mathrm{pol}}\rho_2^{-\nu} + A_{\mathrm{den},S}\{(16+V_S)e^{-Z_S/2} + Q_S\}\right] =: E_W^S. \tag{A.112}$$

The same estimates hold in normalized L^2. For fixed design these imply

$$\left\|f - q_\psi^\star\right\|_\infty \le C_\psi B e^{-c_\psi W_\psi} + 2A_{\mathrm{den},\psi} B Q_\psi(\lambda), \tag{A.113}$$

where $W_G = N + O(\sqrt{N})$, $W_S = \Theta(N)$, $Q_G = O((1+\lambda^{-2})e^{-2\pi^2/\lambda^2})$, and, for fixed ϑ, $Q_S = O_\vartheta(e^{-2\pi\vartheta/\lambda})$ on $0 < \lambda \le 1$.

Proof. Use (A.102), the two quadrature propositions, and (A.110). For $N \ge 2/\eta$, $\nu \ge \eta N/2$. Also $Z_G^2 \ge \lambda^2 a_G^2 N$, $Z_S \ge \lambda a_S N$, and $W_\psi \le (L_*+2)N$. One may therefore take

$$c_G = \frac{\min\{(\eta/2)\log\rho_2, \lambda^2 a_G^2/8\}}{L_*+2}, \qquad c_S = \frac{\min\{(\eta/2)\log\rho_2, \lambda a_S/2\}}{L_*+2}.$$

The stated bounds on Q_ψ follow from their explicit denominators. □

Unlike the tanh correction, the GELU halo estimate uses Gaussian spatial decay directly. The uncorrected SiLU kernel has exponential tails and uses $R_S = \Theta(N)$ to retain a geometric rate in total width.

On the real axis, $|\psi(t)| \le |t|$ and $|\psi'(t)| \le 2$ for both activations: $\psi'_G = \Phi + t\varphi$ has absolute value at most $1+1/\sqrt{2\pi e}$, and $\psi'_S = \varsigma + tg$ has absolute value at most $1+1/e$. Since $(N+2R+1)h \le 2L_*$, write $H_* = 1 + L_*$ and set

$$A_\beta = \delta^{-1} + 4L_* A_{\mathrm{den}}, \qquad A_\alpha = 1 + 2L_* H_* A_{\mathrm{den}} + A_\beta. \tag{A.114}$$

Then (A.93) and (A.108) give

$$|w_j^\star| \le B A_{\mathrm{den}} h/\gamma, \quad |w_\pm^\star| = |\beta|/\gamma \le B A_\beta/\gamma, \quad |\alpha| \le B A_\alpha. \tag{A.115}$$

Here $w_\pm^\star$ are the weights of the two affine units. These inequalities follow respectively from the density bound, $|(q_0^\psi)'(-1)| \le 4L_* A_{\mathrm{den}} B$, and $|q_0^\psi(-1)| \le 2L_* H_* A_{\mathrm{den}} B$.

Let ϕ_j denote each native activation feature, including the affine pair, and $\phi_b = 1$. Choose deterministic positive normalization constants $m_b = 1$ and

$$m_j \ge \max\{1, \sup_{x\in\Omega} |\phi_j(x)|\}, \qquad m_j \le 2(1+\gamma H_*) \tag{A.116}$$

for the grid units, and $m_\pm \le 2(1+\gamma)$ for the affine pair. The exact supremum satisfies these conditions; the least dyadic upper bound on $1+\gamma H_*$ (and $1+\gamma$ for the pair) also does so. These choices may be validated by outward enclosures of the displayed quantities. Write $\varphi_j = \phi_j/m_j$ and $v_j = m_j c_j$. Explicit normalized-coordinate envelopes are

$$\alpha_j = m_j A_{\mathrm{den}} h/\gamma, \quad \alpha_\pm = m_\pm A_\beta/\gamma, \quad \alpha_b = A_\alpha. \tag{A.117}$$

Since $\gamma \ge 1$, their sum satisfies

$$\sum_j \alpha_j \le A_\alpha + 4L_* A_{\mathrm{den}}(H_*+1) + 8A_\beta =: K_{\mathrm{ref},\psi}. \tag{A.118}$$

In particular, this bound includes the bias and both affine units.

Lemma A.4.1 (Normalized features on cell ellipses)**.** *For every fixed $v_0 > 0$ for GELU, and $0 < v_0 < \pi$ for SiLU, there is a constant C_{ψ,v_0} such that*

$$|\psi(u+iv)| \le C_{\psi,v_0}(1+u_+), \qquad |v| \le v_0, \quad u_+ = \max\{u, 0\}. \tag{A.119}$$

Explicit choices are

$$C_{G,v} = 1 + v + \frac{v e^{v^2/2}}{\sqrt{2\pi}}(e^{-1/2} + v), \qquad C_{S,v} = (1+v)\sec(v/2). \tag{A.120}$$

For a core cell of length h and a Bernstein ellipse of parameter $\varrho > 1$, set $t_\varrho = (\varrho + \varrho^{-1} - 2)/4$. If $\lambda(\varrho - \varrho^{-1})/4 \le v_0$, then every normalized feature satisfies

$$|\varphi_j(z)| \le C_{\psi,v_0}(3 + \lambda t_\varrho) \quad \textit{on that ellipse}. \tag{A.121}$$

Proof. For GELU, $|u|\Phi(u)$ is bounded for $u \le 0$ by the Gaussian tail inequality $t\Phi(-t) \le \varphi(t)$. In

$$\psi_G(u+iv) = (u+iv)\Phi(u) + i(u+iv)\int_0^v \varphi(u+it)\,dt,$$

the second term is bounded uniformly in u, because $|\varphi(u+it)| = e^{-u^2/2+t^2/2}/\sqrt{2\pi}$ and $(|u| + v_0)e^{-u^2/2}$ is bounded. This proves (A.119). For SiLU, $|\psi_S(u+iv)| \le \sec(v_0/2)(|u|+v_0)/(1+e^{-u})$; for $u < 0$ use $|u|e^u \le 1/e$, and for $u \ge 0$ use the linear bound.

Each real feature's maximum preactivation on Ω is $U_j = \gamma(1 - x_j)$, or $U_j = \gamma$ for an affine unit. Both activations satisfy $\psi(t) \ge t/2$ for $t \ge 0$. Consequently $m_j \ge \max\{1, (U_j)_+/2\}$. The ellipse extends horizontally by at most ht_ϱ beyond the core; its preactivation real part is at most $U_j + \lambda t_\varrho$. Divide (A.119) by m_j to obtain (A.121). □

Recovery and computed error. Theorem A.4 supplies $E_W^{\mathrm{an}} = E_W^{\psi}$; (A.116)–(A.118) supply the real-axis normalization and deterministic envelopes, including the bias and both affine units. Under (A.55), Lemma A.4.1 supplies $C_{\mathrm{ell}} = C_{\psi,\pi/4}(3 + \lambda t_\varrho)$, with the SiLU ellipses inside its pole-free strip. Thus the common dictionary assumptions in Section A.3 hold. Under that section's numerical, input-rounding, and evaluation conditions, with $M \geq n$, the same computation proves

$$\left\| f - \widehat{q}_W^{\psi} \right\|_{2,\Omega} \leq 4E_W^{\psi} + \epsilon_{\mathrm{rec,W}}^{\psi}, \qquad \left\| \widehat{c}_W^{\psi} \right\|_1 \leq 2BK_{\mathrm{ref},\psi} + e_{c,W}^{\psi}. \tag{A.122}$$

In particular, (A.113) yields geometric refinement in width plus the activation-specific midpoint and numerical allowances. Corollary A.3.1 applies under its fixed-design and certified-range premises.

B SELECTION OF THE RELATIVE BANDWIDTH

We derive the tanh prescription from the pole contribution in (A.52). With the preceding geometry,

$$\frac{2\pi d}{h} = \frac{\pi^2}{\lambda}, \tag{B.1}$$

giving the exponential factor $e^{-\pi^2/\lambda}$ in (A.25).

B.1 FREQUENCY-DEPENDENT POLE BOUND

For a positive angular frequency ω, put

$$\begin{aligned} A_\lambda = \frac{\pi^2}{\lambda}, \qquad t_\omega = \omega d, \qquad \chi_\omega = 2t_\omega = \frac{2\pi\omega}{N\lambda}, \\ \mathcal{T}(\lambda, h\omega) = \sinh t_\omega \left[\frac{1}{\sinh(A_\lambda - t_\omega)} + \frac{1}{\sinh(A_\lambda + t_\omega)}\right], \end{aligned} \tag{B.2}$$

with $\mathcal{T}(\lambda, 0) = 0$. The first-pair score bounds the complete pole contribution as follows.

Proposition B.0.1 (Frequency-dependent aliasing bound)**.** *For a finite exponential sum $f(x) = \sum_j b_j e^{i\omega_j x}$ with $|h\omega_j| < \pi$, the pole contribution on the existing midpoint contour satisfies*

$$E_{\text{alias}} := \sup_{x\in\Omega} |\mathcal{P}_h[F_x]| \le \frac{2}{1 - e^{-A_\lambda}} \sum_{\omega_j \neq 0} |b_j| \mathcal{T}(\lambda, h|\omega_j|). \tag{B.3}$$

Proof. For a unit mode with $\omega > 0$, (A.9) gives $a_\gamma(z) = i\sinh(t_\omega)e^{i\omega z}/d$. At height $(2\ell+1)d$, the density magnitudes at the four poles above and below x, x_0 sum to $4\sinh(t_\omega)\cosh((2\ell+1)t_\omega)/d$. Use the residues $1/(2\gamma)$ in absolute value and (A.21), as in Lemma A.1.3 (Javed & Trefethen, 2014, Secs. 4–5), to obtain

$$\begin{aligned} E_{\text{alias}} &\le 8\sinh t_\omega \sum_{\ell\ge 0} \frac{\cosh((2\ell+1)t_\omega)}{e^{(2\ell+1)A_\lambda} - 1} \\ &= 2\sinh t_\omega \sum_{m\ge 1} \left[\frac{1}{\sinh(mA_\lambda - t_\omega)} + \frac{1}{\sinh(mA_\lambda + t_\omega)}\right]. \end{aligned}$$

Extending the positive residue sum and expanding its geometric denominators gives the second line. Each successive m-term is at most e^{-A_λ} times its predecessor because $\sinh(v + A_\lambda) \ge e^{A_\lambda}\sinh v$ for $v > 0$. Summation proves the unit-mode bound. Negative frequencies have the same absolute bound, constants have zero density, and linearity gives (B.3). The factor two accounts for anchoring at x_0 as well as evaluation at x. □

B.2 COMPACT SCORE AND SCALAR PRESCRIPTION

Define the compact and small-angle scores

$$\mathcal{S}(\lambda, N; \omega) = 2e^{-A_\lambda}\sinh\chi_\omega, \qquad \mathcal{S}_{\text{sa}}(\lambda, N; \omega) = 2\chi_\omega e^{-A_\lambda}. \tag{B.4}$$

For $0 < h\omega < \pi$,

$$1 \le \frac{\mathcal{T}}{\mathcal{S}} \le \frac{1}{1 - e^{-(2A_\lambda - \chi_\omega)}}, \qquad \frac{\mathcal{S}}{\mathcal{S}_{\text{sa}}} = \frac{\sinh\chi_\omega}{\chi_\omega}. \tag{B.5}$$

Indeed, expand $1/\sinh z = 2e^{-z}/(1 - e^{-2z})$ in the two terms of (B.2). Their leading exponentials sum to $4e^{-A_\lambda}\sinh t_\omega \cosh t_\omega = \mathcal{S}$, and the smaller denominator is $1 - e^{-(2A_\lambda - \chi_\omega)}$. The compact approximation is therefore controlled throughout this band; small-angle linearization additionally requires $\chi_\omega \ll 1$. Figure 4 shows the width dependence of the refined bandwidth selections.

For a representative frequency $\bar\omega$ and a dimensionless score budget $\varepsilon_{\text{eff}} > 0$, the practical rule solves

$$\boxed{2\exp\left(-\frac{\pi^2}{\lambda_{\text{pred}}}\right)\sinh\left(\frac{2\pi\bar\omega}{N\lambda_{\text{pred}}}\right) = \varepsilon_{\text{eff}}.} \tag{B.6}$$

Use an admissible bandwidth bracket on an increasing branch. In fact,

$$\frac{\partial \log \mathcal{S}}{\partial \log \lambda} = A_\lambda - \chi_\omega \coth \chi_\omega > 0 \quad \text{if } A_\lambda - \chi_\omega \geq 1, \tag{B.7}$$

because $\chi \coth \chi < \chi + 1$. Check the condition at the largest bracket endpoint and use bisection with $\log \mathcal{S} = -(A_\lambda - \chi_\omega) + \log(-\operatorname{expm1}(-2\chi_\omega))$. If the whole bracket satisfies the budget, select its largest endpoint; if none does, report infeasibility. A zero frequency imposes no score constraint.

B.3 Frequency input and tolerance

For known coefficients use $\bar{\omega} = \sum_j |b_j||\omega_j| / \sum_j |b_j|$. For period-two sample data, take $m_F = M - 1 = sN$ distinct observations $y_\ell = f(-1 + 2\ell/m_F)$, excluding the repeated endpoint from the transform, and define

$$\widehat{b}_k = \frac{1}{m_F} \sum_{\ell=0}^{m_F-1} y_\ell e^{-2\pi i k \ell / m_F}, \qquad \widehat{\bar{\omega}} = \frac{\sum_{k \in \mathcal{K}} |\widehat{b}_k|\, |\pi k|}{\sum_{k \in \mathcal{K}} |\widehat{b}_k|}, \qquad \mathcal{K} = \{-\lfloor m_F/2 \rfloor, \dots, \lceil m_F/2 \rceil - 1\}. \tag{B.8}$$

Weights are amplitudes, not powers; the constant mode is included in the denominator. A zero denominator returns no estimate. For nonperiodic targets, the endpoint treatment and any spectral filtering must be specified: the raw periodic transform can reflect endpoint mismatch rather than interior variation. m_F controls observation resolution; N controls hidden-grid replicas.

Amplitude weighting matches the leading small-angle bound, which is proportional to $\sum_j |b_j||\omega_j|$. At finite angles, however, convexity of $\sinh$ gives $\big(\sum_j |b_j|\big)\mathcal{S}(\bar{\omega}) \leq \sum_j |b_j| \mathcal{S}(|\omega_j|)$, with λ, N held fixed. The centroid substitution is therefore a heuristic, not an upper-bound-preserving reduction. To certify aliasing, use (B.3) with all relevant frequencies, or the spectrum-free bound (A.25).

For one nonzero absolute frequency with coefficient mass $B_{\text{spec}} = \sum_j |b_j|$, an absolute alias budget ε_a gives the leading normalization $\varepsilon_{\text{eff}} \simeq \varepsilon_a / (2B_{\text{spec}})$. A bound retains the corrections in (B.3) and (B.5); a relative budget first multiplies by $\|f\|_{2,\Omega}$. Allocate the budget alongside the other allowances in Section A.3, rather than necessarily setting it to unit roundoff. The prescription is an aliasing criterion, not a characterization of the total-error minimizer.

Width dependence and induced slopes. Along an interior root of (B.6), with frequency and tolerance fixed, implicit differentiation in continuous N gives

$$\frac{d \log \lambda_{\text{pred}}}{d \log N} = \frac{\chi \coth \chi}{A_\lambda - \chi \coth \chi}, \qquad \frac{\gamma_{\text{pred}}}{W} = \frac{\lambda_{\text{pred}}}{2} \frac{N}{N + 2\lceil a_H \sqrt{N} \rceil + 1}, \tag{B.9}$$

where $\omega = \bar{\omega}$ and $\lambda = \lambda_{\text{pred}}$ in the definitions above. Thus the relative bandwidth varies slowly when $A_\lambda \gg \chi \coth \chi$. Fixed λ gives $\gamma / W \to \lambda / 2$; more generally, uniform positive lower and upper bounds on the selected bandwidth give $\gamma_{\text{pred}} = \Theta(W)$ on that regime. Figure 4 shows the finite-width behavior, without asserting a width-independent asymptotic limit. Figure 5 compares the selections with measured approximation errors.

B.4 Activation-specific rules used in the figures

For an activation σ, let $K = \sigma^{(r)}$ be its localized derivative, with $K \in L^1(\mathbb{R})$ and $\widehat{K}(0) \neq 0$. Use $\widehat{K}(\xi) = \int K(t) e^{-i\xi t}\, dt$ and define the general rule by

$$\frac{|\widehat{K}(2\pi/\lambda_{\text{general}})|}{|\widehat{K}(0)|} = \varepsilon_{\text{eff}}. \tag{B.10}$$

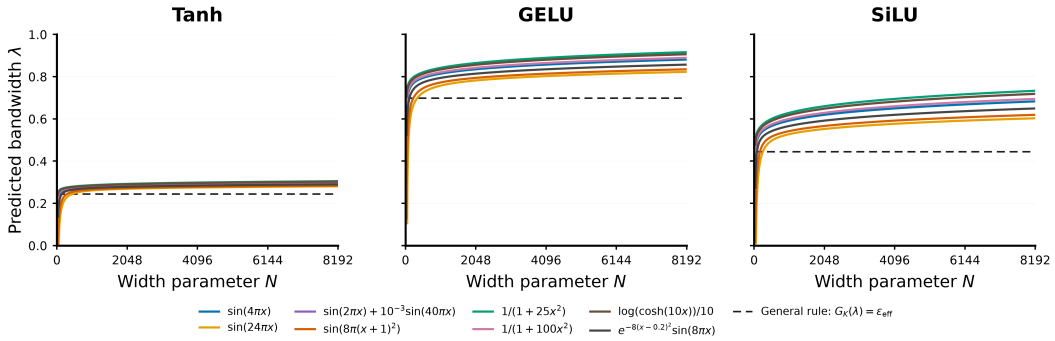


Figure 4: Bandwidth selection versus width parameter N for tanh, exact GELU, and SiLU. Colored curves give the refined selections for the eight targets in the legend, using $\varepsilon_{\text{eff}} = 2^{-52}$. Dashed lines give the activation-specific general-rule selections (approximately 0.25 for tanh). These are scalar rule calculations, not fitted-network errors; the refined selections retain their frequency-dependent drift.

The transforms used in the figures are

$\sigma(x)$	r	$\widehat{K}(\xi)$
$\tanh x$	1	$\dfrac{\pi\xi}{\sinh(\pi\xi/2)}$
$x\Phi(x)$ (exact GELU)	2	$(1+\xi^2)e^{-\xi^2/2}$
$x/(1+e^{-x})$ (SiLU)	2	$\dfrac{(\pi\xi)^2\cosh(\pi\xi)}{\sinh^2(\pi\xi)}$
$\operatorname{erf}(x)$	1	$2e^{-\xi^2/4}$

(B.11)

Here Φ is the standard normal cumulative distribution function; the transform values at zero are $2, 1, 1, 2$, respectively. For tanh and $\varepsilon_{\text{eff}} = 2^{-52}$, equation B.10 gives $\lambda_{\text{general}} \simeq 0.2441$; $\lambda = 0.25$ is its rounded practical default.

For $\theta = h\bar{\omega} = 2\bar{\omega}/N$, the refined markers solve

$$T_{K,r}(\lambda,\theta) := \sum_{\ell\in\{-1,1\}} \left|\frac{\theta}{\theta+2\pi\ell}\right|^r \frac{|\widehat{K}((\theta+2\pi\ell)/\lambda)|}{|\widehat{K}(\theta/\lambda)|} = \varepsilon_{\text{eff}}, \qquad 0<\theta<\pi. \tag{B.12}$$

This is the magnitude sum of the nearest Fourier replicas of a central-amplitude-matched infinite lattice. Integrating the rth derivative back to the function supplies the frequency factor raised to r. For tanh it is exactly (B.2); the main-text rule uses its controlled compact approximation (B.6). Select a root on an admissible increasing branch with nonzero denominator. The other activations use this replica score as a bandwidth heuristic; the finite-interval pole-contribution bound above is specific to tanh. Replacing a spectrum by $\bar{\omega}$ retains the limitations already stated.

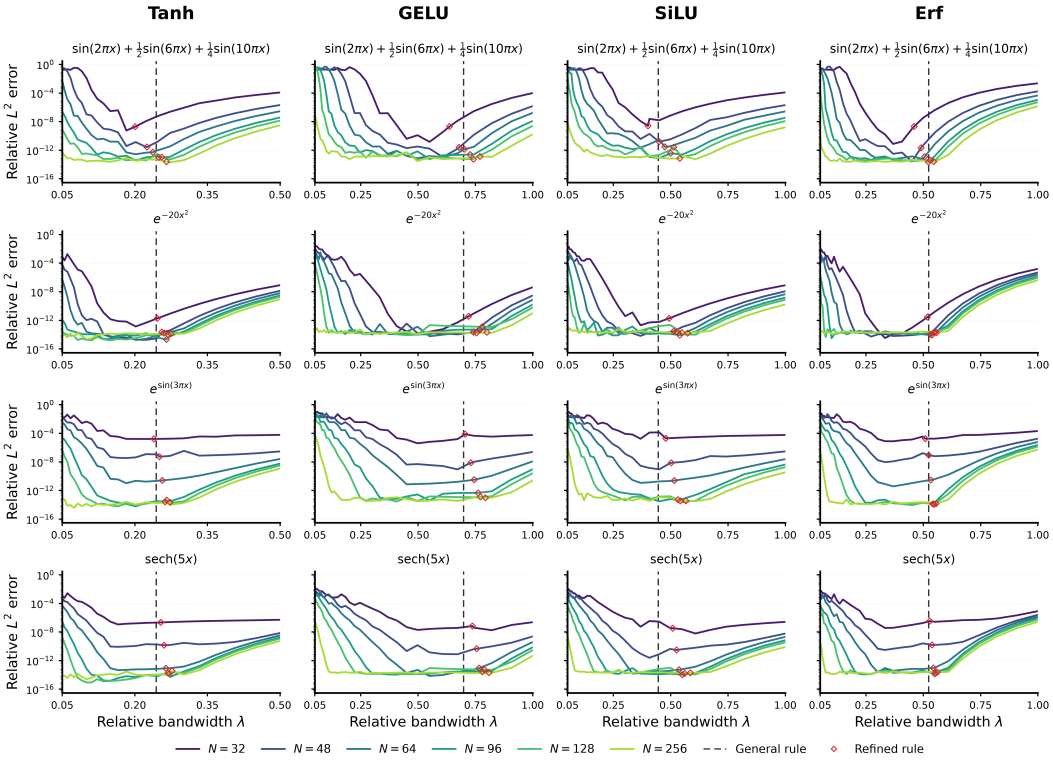


Figure 5: Measured relative L^2 error versus bandwidth for four targets (rows) and four activations (columns). Colors indicate $N = 32, 48, 64, 96, 128, 256$, with 32 halo nodes per side and total width $W = N + 65$. Vertical dashed lines mark the general rule; red diamonds mark the refined selections and their directly measured errors. Readouts use FP64 SVD least squares on 2003 uniform points, with relative singular-value cutoff 2^{-52}; errors use 4001 uniform evaluation points. Tanh is shown over $0.05 \le \lambda \le 0.5$; the other activations extend to 1. Refined selections use the representative-frequency heuristic and need not minimize measured error.

C PRECISION-CONTROLLED COMPUTATION

The precision experiment uses a binary format with p significand bits, including the leading bit, and rounds every elementary operation during feature construction, readout recovery, and evaluation.

Arithmetic. Let fl_p denote rounding to the chosen format using round-to-nearest, ties-to-even. Arithmetic primitives return

$$\mathrm{fl}_p(a+b), \quad \mathrm{fl}_p(a-b), \quad \mathrm{fl}_p(ab), \quad \mathrm{fl}_p(a/b), \quad \mathrm{fl}_p(\sqrt{a}).$$

Operands and returned values belong to the working format. A model multiplication and its subsequent addition are rounded separately, and dot products use sequential rounded accumulation. Constants are prepared beforehand and rounded into the format. The unit roundoff for normal results is $u_p = 2^{-p}$.

The sweep uses $p = 8, \ldots, 53$, with normal exponents ranging from -958 to 959 and gradual underflow. This expanded exponent range isolates the effect of significand precision; in particular, the $p = 24$ sweep point is not a native binary32 computation. Separate checks use the IEEE binary32 and binary64 exponent ranges. The emulator stores values in binary64 containers, but each arithmetic primitive returns a value in the selected format. Its rounding implementation resolves midpoint cases to avoid double-rounding errors.

Feature formation and evaluation. Training locations, target samples, and the selected bandwidth are rounded before entering the computation. Grid spacing, centers, and slopes are then formed using the rounded arithmetic. Features use the centered evaluation order

$$\phi_{j,p}(x) = \tanh_p(\mathrm{fl}_p(\gamma\ \mathrm{fl}_p(x - c_j))),$$

where x, c_j, and γ are working-format values. The activation $\tanh_p$ uses range reduction and polynomial evaluation built from the same rounded primitives. Near the origin, an exponential-difference formulation avoids subtracting nearly equal numbers. It is a specified numerical approximation to tanh, rather than an assumption of a correctly rounded transcendental operation.

The design matrix contains these features and a constant bias column. After fitting, coefficients remain in the working format. Prediction starts from the fitted bias and adds feature contributions sequentially in increasing center order, rounding each multiplication and addition. At low precision, distinct centers or sample locations can round to the same value; this effect is retained in the experiment.

Least-squares solver. We port reference LAPACK 3.12.1's `DGELSS` SVD least-squares driver operation by operation, using its unblocked configuration for an overdetermined system with one right-hand side. The computation includes Householder QR reduction, bidiagonalization, iterative diagonalization of the bidiagonal matrix, and truncated back-substitution. Norms, rotations, matrix–vector products, and scaling operations use the working arithmetic. Machine constants, including precision and safe-minimum thresholds, are set for that format.

Let $\widehat{\sigma}_1 \geq \cdots \geq \widehat{\sigma}_{W+1}$ denote the computed singular values. We retain precisely those satisfying

$$\widehat{\sigma}_i > \max\{\mathrm{fl}_p(2^{1-p}\widehat{\sigma}_1), \mathrm{sfmin}_p\},$$

where sfmin_p is LAPACK's safe minimum for the chosen format. Thus the primary relative cutoff is $\texttt{RCOND} = 2^{1-p} = 2u_p$, fixed before observing the error. The empirical solver uses the unscaled feature matrix; the reference-scaled recovery theorem specifies a separate algorithm and additional numerical conditions.

Experimental protocol. The precision panel uses the chirp $f(x) = \sin(8\pi(x+1)^2)$ on $[-1, 1]$, total width $W = 1024$, 975 interior intervals, and 24 halo centers per side. Training and evaluation use 4,801 and 8,001 equispaced points, respectively. For each p, the refined bandwidth rule is evaluated offline with tolerance 2^{1-p}, and its result is rounded on entry. Target sampling also occurs outside the controlled computation.

Reported relative errors are measured afterwards in binary64 against the target at the original, unrounded evaluation locations. Consequently, input quantization contributes to the measured error. The experiment measures finite-grid relative error; the external error calculation does not supply additional precision to the fitted model.

Validation. Arithmetic primitives were checked against MPFR, including subnormal and midpoint cases. An independent MPFR replay verifies the operation sequence for activation evaluation, feature formation, and readout summation. Native-format solver checks reproduce reference LAPACK `SGELSS`/`DGELSS` bit for bit on the tested problems, including the full 4801×1025 feature matrix. A static audit checks that numerical operations pass through the rounded primitives. Dimension-based solver branch selection remains outside the controlled numerical kernels. These checks validate the implemented computation; the observed precision–error relation is an empirical result.

D Derivation of the comparison bounds

We derive the implementation-specific entries in Table 1 for the fixed analytic class of Section A.1. Write $L_\varepsilon = \log(1/\varepsilon)$, with $0 < \varepsilon < e^{-2}$, and use W for maximum hidden-layer width, $P_{\max}$ for the largest absolute weight or bias, and b_{work} for working significand precision. Constants depend on the fixed class and construction, not on ε. The entries are sufficient bounds for the realizations specified below.

Common approximation and arithmetic estimates. Use round-to-nearest arithmetic with $u = 2^{-b_{\mathrm{work}}}$, accurately rounded elementary functions, and sufficient exponent range. Each construction requests its own sample locations, with absolute sample error $O(u)$; location and input rounding are included using the class's uniform Lipschitz bound. The standard operationwise and summation estimates in (Higham, 2002, Chs. 2–4) are used throughout, with their smallness conditions satisfied by the precisions below. If coefficient formation, storage, and evaluation contribute at most $u\mathcal{A}$, an arithmetic budget of $\varepsilon/2$ requires only

$$b_{\mathrm{work}} \geq \log_2(1/\varepsilon) + \log_2 \max\{1, \mathcal{A}\} + O(1). \tag{D.1}$$

For the polynomial-based rows, the degree-ν Chebyshev interpolant at $\cos(\ell\pi/\nu)$, $0 \leq \ell \leq \nu$, satisfies

$$P_\nu = \sum_{j=0}^{\nu} c_{j,\nu} T_j, \qquad \|f - P_\nu\|_\infty \leq C\rho^{-\nu}, \qquad |c_{j,\nu}| \leq C\rho^{-j}, \quad \rho > 1, \tag{D.2}$$

by (Trefethen, 2019, Chs. 4 and 8). Thus $\nu = O(L_\varepsilon)$ suffices and $\sum_j (1+j)|c_{j,\nu}| \leq C$. Computing the coefficients from the samples by the discrete cosine formula gives $\|\widehat{\boldsymbol{c}} - \boldsymbol{c}\|_1 \leq u\ \mathrm{poly}(\nu)$ by the same arithmetic estimates, even with direct summation. Here $\mathrm{poly}(\nu)$ denotes a fixed-degree polynomial bound. The competitor function-error bounds below are uniform and therefore also imply normalized RMS bounds.

D.1 Classical staircase

Let $x_j = -1 + 2j/N$ and $S(t) = (1 + \tanh t)/2$. The specified staircase is

$$q_N(x) = f(-1) + \sum_{j=1}^{N} \Delta_j f\, S(N(x - x_j)), \qquad \Delta_j f = f(x_j) - f(x_{j-1}). \tag{D.3}$$

The historical reference is Jones (1990); the rates here concern (D.3). If L_f is the Lipschitz constant, then $|\Delta_j f| \leq 2L_f/N$. Replacing S by the Heaviside step gives error at most $2L_f/N$, while $|S(t) - H(t)| \leq e^{-2|t|}$ gives

$$|q_N - q_N^{\mathrm{step}}| \leq \frac{2L_f}{N} \sum_{j=1}^{N} e^{-2N|x - x_j|} \leq \frac{CL_f}{N}.$$

The last sum is bounded uniformly because the scaled centers are separated by two. Expanding S yields N tanh units, hidden parameters $O(N)$, readout weights $\Delta_j f/2$ with absolute sum at most L_f, and bias $[f(-1) + f(1)]/2$. Sample errors accumulate to $O(uN)$; affine errors $O(uN)$ per feature and sequential summation have the same output bound because the readout sum is bounded (Higham, 2002, Chs. 3–4). Hence

$$\|f - \widehat{q}_N\|_\infty \leq C/N + CuN. \tag{D.4}$$

Taking $N \asymp \varepsilon^{-1}$ and $u \lesssim \varepsilon/N$ gives $W, P_{\max} = O(\varepsilon^{-1})$, $b_{\mathrm{work}} = O(L_\varepsilon)$, and one hidden layer.

D.2 A polynomial-based tanh specialization

For the row associated with Mhaskar (1996), consider the following explicit realization of the polynomial-approximation strategy. Write $P_\nu = \sum_{p=0}^{\nu} \mu_p x^p$. The Chebyshev recurrence gives $\sum_p |\mu_p| \leq Ce^{C\nu}$, since the absolute monomial coefficient sum of T_j is at most 3^j. Set $\beta_* = \frac{1}{2}\log 2$ and, for dyadic $\eta > 0$, define

$$H_{p,\eta}(x) = \frac{\sum_{r=0}^{p} (-1)^{p-r} \binom{p}{r} \tanh\big(\beta_* + (r - p/2)\eta x\big)}{\eta^p \tanh^{(p)}(\beta_*)}, \quad q_{\nu,\eta} = \mu_0 + \sum_{p=1}^{\nu} \mu_p H_{p,\eta}. \tag{D.5}$$

All features have the form $\tanh(\beta_* + j\eta x/2)$, $-\nu \le j \le \nu$; combining repeated features uses at most $2\nu + 1$ units.

The denominators admit an explicit lower bound. Define $E_1(z) = 1$ and $E_{p+1}(z) = (1 + pz)E_p(z) + z(1-z)E_p'(z)$. Differentiation gives

$$\tanh^{(p)}(\beta_*) = \frac{2^{p+2}E_p(-2)}{3^{p+1}}, \qquad |\tanh^{(p)}(\beta_*)|^{-1} \le \frac{3^{p+1}}{2^{p+2}}. \tag{D.6}$$

Indeed, E_p has integer coefficients and constant coefficient one, so $E_p(-2)$ is odd and nonzero. Repeated integration of the centered difference expresses $H_{p,\eta}$ as x^p times the average of $\tanh^{(p)}(\beta_* + \eta x \sum_i t_i)/\tanh^{(p)}(\beta_*)$ on $[-1/2, 1/2]^p$. The sum has mean zero and second moment $p/12$, giving

$$\|H_{p,\eta} - x^p\|_\infty \le \frac{p\eta^2 \|\tanh^{(p+2)}\|_\infty}{24|\tanh^{(p)}(\beta_*)|}. \tag{D.7}$$

Cauchy's estimate on a fixed narrower tanh strip bounds the derivative by $(p+2)!C^{p+2}$. Consequently $\|q_{\nu,\eta} - P_\nu\|_\infty \le \eta^2 e^{C\nu\log(\nu+2)}$. Choose $\nu = O(L_\varepsilon)$ and dyadic η, with $\nu\eta \le 1$, so both approximation errors are at most $\varepsilon/4$. This permits

$$\log(1/\eta) = O\big(L_\varepsilon + \nu\log(\nu+2)\big), \qquad |\mu_0| + \sum_j |w_j| \le Ce^{C\nu}(2/\eta)^\nu. \tag{D.8}$$

The second bound follows by summing the uncombined binomial coefficients and using (D.6); it also bounds the combined readout. Hidden slopes are at most $1/2$ and hidden biases equal β_*.

For coefficient formation, compute the sampled Chebyshev coefficients, convert by the polynomial recurrence, and accumulate the shared-feature weights. Conversion has amplification $\mathrm{poly}(\nu)e^{C\nu}$. The recurrence for E_p has nonnegative coefficients summing to $p!$, so $|E_p(-2)| \le 2^{p-1}p!$: its integer intermediates need only $O(\nu\log(\nu+2))$ bits, within the budget below. Powers of dyadic η are exponent shifts. Rounded divisions, storage of β_*, and bounded feature arguments then give, by (Higham, 2002, Chs. 2–4),

$$\|\widehat{q}_{\nu,\eta} - q_{\nu,\eta}\|_\infty \le u\, e^{C\nu\log(\nu+2)}(2/\eta)^\nu. \tag{D.9}$$

This is an absolute error bound and does not assume relative accuracy of cancelling sums. Substituting (D.8) in (D.1) gives one hidden layer, $W = O(L_\varepsilon)$, $P_{\max} \le \exp[O(L_\varepsilon^2 \log L_\varepsilon)]$, and $b_{\mathrm{work}} = O(L_\varepsilon^2 \log L_\varepsilon) = \widetilde{O}(L_\varepsilon^2)$.

D.3 Sampled sigmoidal quasi-interpolation

Use $\phi = S'$ and $\psi(t) = S(t+1) - S(t)$ with the same S as above. The translates of ψ form a nonnegative partition of unity (Costarelli & Spigler, 2015, Lemma 5.1). Extend f to a bounded Lipschitz function F by linear interpolation from its endpoint values to zero at ± 2, and set $F = 0$ outside $[-2, 2]$. The exponential kernel tails satisfy the hypotheses of (Costarelli & Spigler, 2015, Theorem 5.4(ii)); with scale N and cutoff $2N$, this gives

$$\left\| f - \sum_{k=-2N}^{2N} a_k \psi(N\,\cdot - k) \right\|_\infty \le C/N, \qquad a_k = \int_{\mathbb{R}} \phi(t) F((k+t)/N)\, dt. \tag{D.10}$$

Replace the coefficients by samples. Since $\int \phi = 1$,

$$|a_k - F(k/N)| \le \frac{L_F}{N} \int_{\mathbb{R}} |t|\phi(t)\, dt \le C/N.$$

The partition of unity bounds the function change by the largest coefficient change, so $q_N^{\mathrm{CS}} = \sum_{k=-2N}^{2N} F(k/N)\psi(N\,\cdot - k)$ retains error C/N. The extension uses only core samples and endpoint values. Expanding each ψ into two tanhs without merging gives at most $8N + 2$ units, hidden parameters $O(N)$, and absolute readout sum $O(N)$. Affine errors $O(uN)$ and standard summation bounds therefore give

$$\|f - \widehat{q}_N^{\mathrm{CS}}\|_\infty \le C/N + CuN^2 \tag{D.11}$$

including coefficient formation (Higham, 2002, Chs. 3–4). Choose $N \asymp \varepsilon^{-1}$ and $u \lesssim \varepsilon/N^2$ to obtain $W, P_{\max} = O(\varepsilon^{-1})$, $b_{\mathrm{work}} = O(L_\varepsilon)$, and one hidden layer.

D.4 CHEBNET

For the ReQU activation $\sigma_2(t) = (\max\{t, 0\})^2$, the recursive realization in (Tang et al., 2019, Sec. 2.1, Theorem 1, 2019 version) represents P_ν exactly with $O(\nu)$ neurons and nonzero weights and $O(\log(\nu + 1))$ hidden layers. Its recursive split uses $T_{m+j} = 2T_m T_j - T_{m-j}$, with m a power of two. Along a path into a fixed subtree, a doubling sends index $m + j$ to $j < (m + j)/2$, while a reflection changes only the sign. Each original coefficient has at most one path into that subtree. Its amplification is therefore at most $2\max\{1, j\}$, and the intermediate coefficient sums obey

$$\sum_k |c_k^{\text{subtree}}| \le 2 \sum_{j=0}^{\nu} \max\{1, j\}|c_{j,\nu}| \le C. \tag{D.12}$$

Thus intermediate polynomials are bounded on $[-1, 1]$. The fixed-size square, product, and identity modules of (Tang et al., 2019, Lemma 2) then have uniformly bounded weights, biases, and exact hidden states; in particular, $P_{\max} = O(1)$.

The coefficient recursion involves additions and doublings through logarithmic depth, so parameter formation has error $u\ \text{poly}(\nu)$. On a fixed enlargement of the exact state ranges, each bounded-fan-in module is uniformly Lipschitz and has local rounding error $O(u)$ (Higham, 2002, Chs. 2–3). Hence the maximum layer error satisfies

$$E_{\ell+1} \le C_1 E_\ell + u\ \text{poly}(\nu), \qquad E_0 = O(u), \qquad \ell = O(\log(\nu + 1)).$$

It follows that $\|\widehat{q}_\nu - P_\nu\|_\infty \le u\ \text{poly}(\nu)$; choosing this bound below $\varepsilon/2$ also keeps all layer errors below one, closing the enlarged-range induction. Together with (D.2), this gives $W = O(L_\varepsilon)$, depth $O(\log L_\varepsilon)$, and $b_{\text{work}} = \log_2(1/\varepsilon) + O(\log L_\varepsilon)$.

D.5 QUILL

Theorem A.1 and Corollary A.3.1 supply this row in their geometric-to-allowance regime. With the allowance at most $\varepsilon/2$ and the resulting width in the certified range, $W = O(L_\varepsilon)$ suffices. Bounded centers and readout, together with $\gamma = \lambda/h = O(W)$, give $P_{\max} = O(W)$ and one hidden layer. For a construction-and-evaluation allowance $uC_{\text{num}}W^q$ uniform on that regime, (D.1) gives $b_{\text{work}} = \log_2(1/\varepsilon) + O(\log W) = O(L_\varepsilon)$. These substitutions retain the preceding aliasing, sampling, and numerical conditions; varying the design with ε additionally requires uniform constants along that choice.

D.6 MEASURED COMPARISON AT FIXED PRECISION

Figure 6 complements Table 1 with measured error versus total hidden-neuron budget at two working precisions. QUILL reaches small errors with fewer neurons in these comparisons; ChebNet attains lower final errors at larger budgets.

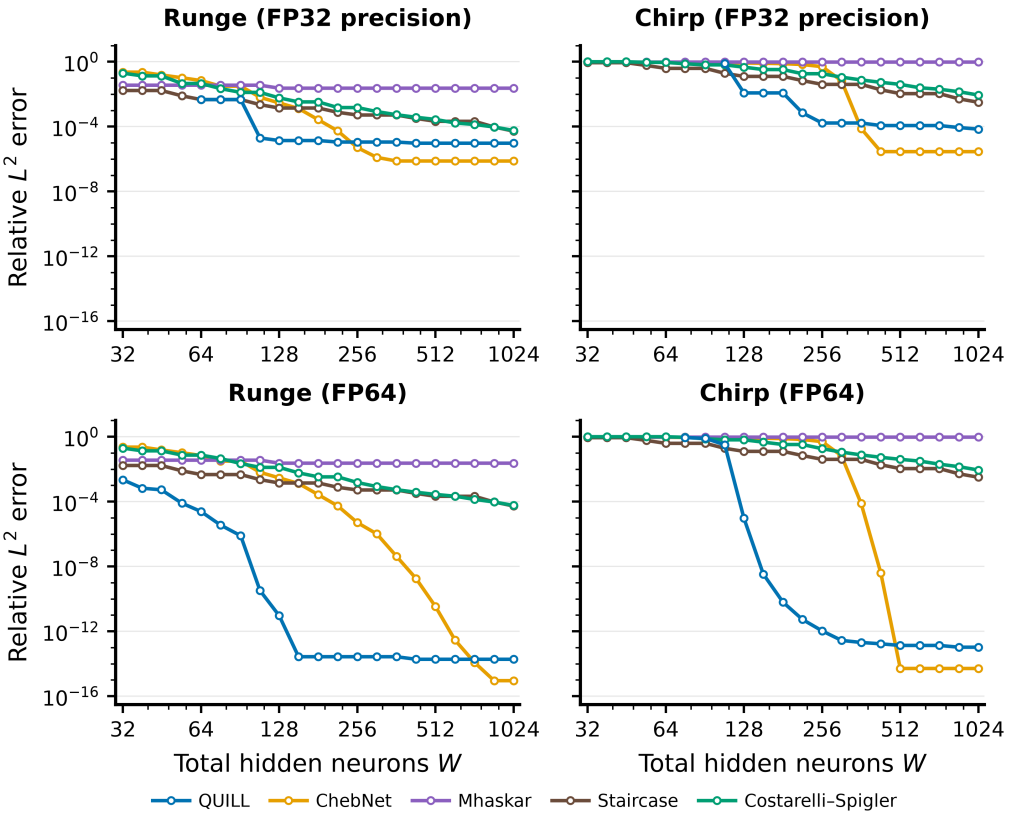


Figure 6: Measured accuracy versus total hidden-neuron budget for Runge, $f(x) = 1/(1 + 25x^2)$ (left), and the chirp, $f(x) = \sin(8\pi(x + 1)^2)$ (right), on $[-1, 1]$. Each marker gives the smallest validation relative L^2 error among saved candidates with at most W hidden neurons, including all layers and halo neurons; plateaus may retain a smaller candidate. Top: 24-bit significands with an expanded exponent range (FP32 precision, rather than native binary32 arithmetic). Bottom: native FP64 for QUILL and ChebNet, and saved 53-bit arithmetic runs for the other methods. QUILL's 24-bit candidate grid is coarser; budgets without a saved candidate are left blank. All panels share the same scales.

E FINITE-KERNEL RATE BOUNDS AND PROOF OF THEOREM 3.2

We bound the normalized eigenvalues of the actual finite tanh dictionary. The continuous-center calculation makes the slope dependence explicit; the lattice and exterior-center corrections transfer it to the finite network. All bounds below are evaluated at the given γ, without a reference slope.

E.1 TRAINING PROBLEM AND EXACT RESIDUAL FORMULA.

Let $z_1, \dots, z_M \in [-1, 1]$, $M \geq 2$, be the training inputs and $x_1, \dots, x_W$ consecutive centers on a lattice $x_{\mathrm{lat}} + h\mathbb{Z}$, with $h > 0$ and $\gamma > 0$. Define the normalized feature matrix and target by

$$(A_\gamma)_{a0} = M^{-1/2}, \qquad (A_\gamma)_{aj} = M^{-1/2} \tanh(\gamma(z_a - x_j)), \qquad y_a = M^{-1/2} f(z_a), \qquad \Phi_\gamma = A_\gamma A_\gamma^T. \tag{E.1}$$

Thus $L(c) = \frac{1}{2}\|A_\gamma c - y\|_2^2$ is the sampled squared-error objective with factor $1/(2M)$. For $r_n = A_\gamma c_n - y$,

$$c_{n+1} = c_n - \eta A_\gamma^T r_n, \qquad r_{n+1} = (I - \eta \Phi_\gamma) r_n, \qquad c_0 = 0. \tag{E.2}$$

Let $(\mu_i, u_i)_{i=1}^M$ be orthonormal eigenpairs of Φ_γ, ordered by decreasing eigenvalue. The bias implies $\mu_1 \geq 1$. For $y \neq 0$, set

$$\rho_i = \mu_i/\mu_1, \qquad p_i = |u_i^T y|^2/\|y\|_2^2, \qquad \eta = 1/(2\mu_1). \tag{E.3}$$

Diagonalizing the recurrence gives

$$E_\gamma(n)^2 := \frac{\|r_n\|_2^2}{\|y\|_2^2} = \sum_{i=1}^M p_i (1 - \rho_i/2)^{2n}. \tag{E.4}$$

This equals the relative squared training RMS error in the main theorem. Zero eigenvalues are included; their target components do not decay.

E.2 CONTINUOUS-CENTER MATRIX AND THE FOURIER MULTIPLIER.

Write $t_\gamma(c) = (\tanh(\gamma(z_a - c)))_{a=1}^M$, and choose $Q \in \mathbb{R}^{M \times (M-1)}$ with $Q^T Q = I$ and $Q^T \mathbf{1} = 0$. The projection removes the constant tails, making the following integral finite:

$$H_\gamma = \frac{1}{Mh} \int_{\mathbb{R}} Q^T t_\gamma(c) t_\gamma(c)^T Q \, dc = -\frac{2}{Mh} Q^T [d_{ab} \coth(\gamma d_{ab})]_{a,b} Q, \quad d_{ab} = z_a - z_b. \tag{E.5}$$

At $d_{ab} = 0$ the bracket is interpreted as $1/\gamma$. Indeed, $1 - \tanh u \tanh v = \coth(u - v)(\tanh u - \tanh v)$ gives $\int_{\mathbb{R}} [1 - t_{\gamma,a}(c) t_{\gamma,b}(c)] \, dc = 2 d_{ab} \coth(\gamma d_{ab})$; the constant term vanishes under Q. No eigenvector of Φ_γ is assumed to have zero mean.

With Fourier convention $\widehat{g}(\omega) = \int g(x) e^{-i\omega x} \, dx$, let $\varrho_\gamma(x) = \gamma \operatorname{sech}^2(\gamma x)/2$. Then

$$\tanh(\gamma x) = (\varrho_\gamma * \operatorname{sgn})(x), \qquad M_\gamma(\omega) = \widehat{\varrho}_\gamma(\omega) = \frac{z}{\sinh z}, \quad z = \frac{\pi|\omega|}{2\gamma}, \quad M_\gamma(0) = 1. \tag{E.6}$$

The convolution follows by integrating ϱ_γ up to x. For the transform, substitution $s = e^{2\gamma x}$ gives $\int_0^\infty s^{-i\nu}(1+s)^{-2} \, ds = \Gamma(1 - i\nu)\Gamma(1 + i\nu) = \pi\nu/\sinh(\pi\nu)$, where $\nu = \omega/(2\gamma)$. For $u = Qv$, the step sum $\sum_a u_a \operatorname{sgn}(c - z_a)$ is compactly supported and has derivative $2\sum_a u_a \delta_{z_a}$. Convolution followed by Parseval therefore yields

$$v^T H_\gamma v = \frac{2}{\pi M h} \int_{\mathbb{R}} \frac{M_\gamma(\omega)^2}{\omega^2} \left| \sum_a (Qv)_a e^{-i\omega z_a} \right|^2 d\omega. \tag{E.7}$$

The apparent singularity at zero is removable because $\mathbf{1}^T Q v = 0$. Every integrand contribution is nonnegative, and its only slope dependence is M_γ^2. In particular,

$$\frac{M_{2\gamma}(\omega)^2}{M_\gamma(\omega)^2} = \cosh^2\left(\frac{\pi|\omega|}{4\gamma}\right). \tag{E.8}$$

This factor is close to one for $|\omega| \ll \gamma$ and exponentially large for $|\omega| \gg \gamma$. Thus contributions at different frequencies change by different factors. Since M_γ increases with γ, $H_{\gamma_b} \succeq H_{\gamma_a}$ for $\gamma_b \geq \gamma_a$; the minimum–maximum principle gives the same ordering of their eigenvalues. The finite-network normalization and corrections must still be retained.

E.3 Explicit lattice and exterior-center corrections.

For $\nu \neq 0$ the integrable product deficit $F_{ab}(c) = t_{\gamma,a}(c)t_{\gamma,b}(c) - 1$ has transform

$$\widehat{F}_{ab}(\nu) = -\frac{4M_\gamma(\nu)}{\nu} \sin(\nu d_{ab}/2) \coth(\gamma d_{ab}) e^{-i\nu(z_a+z_b)/2}. \tag{E.9}$$

At $d_{ab} = 0$ the limit is $-2M_\gamma(\nu)e^{-i\nu z_a}/\gamma$. This follows by transforming the integrable difference of tanhs in the identity used for equation E.5. Choose an integer $J \geq 0$, put $\nu_k = 2\pi k/h$, and retain

$$G_\gamma = \frac{2}{Mh} \sum_{k=1}^{J} Q^T \operatorname{Re}\Big[e^{i\nu_k x_{\mathrm{lat}}} \widehat{F}(\nu_k)\Big] Q. \tag{E.10}$$

Let $\mathcal{F}$ be a finite set of lattice centers absent from the network. Choose it so the remaining absent centers are precisely two exterior tails, starting at $c_R > z_{\max}$ and $c_L < z_{\min}$; in particular, include any absent centers within the sample interval. Define

$$T_\gamma = \frac{1}{M} \sum_{c \in \mathcal{F}} Q^T t_\gamma(c) t_\gamma(c)^T Q, \qquad S_\gamma = H_\gamma + G_\gamma - T_\gamma. \tag{E.11}$$

For any $0 < \vartheta < \pi/2$, set $\bar{z} = M^{-1} \sum_a z_a$, $d_R = c_R - z_{\max} > 0$, and $d_L = z_{\min} - c_L > 0$. The unretained terms obey

$$\delta_\gamma = \frac{8\gamma \|z - \bar{z}\mathbf{1}\|_2^2 \sec^4 \vartheta}{3Mh} \frac{e^{-2\pi\vartheta(J+1)/(\gamma h)}}{1 - e^{-2\pi\vartheta/(\gamma h)}}, \tag{E.12}$$

$$\tau_\gamma = \frac{4(e^{-4\gamma d_R} + e^{-4\gamma d_L})}{1 - e^{-4\gamma h}}, \tag{E.13}$$

$$S_\gamma - (\delta_\gamma + \tau_\gamma) I \preceq Q^T \Phi_\gamma Q \preceq S_\gamma + \delta_\gamma I. \tag{E.14}$$

Here $z = (z_a)_{a=1}^M$. The retained matrices G_γ, T_γ need not be small; the allowances bound only their unretained terms.

To prove the lattice allowance, let $u = Qv$ and $g(c) = \sum_a u_a \tanh(\gamma(c - z_a))$. Subtracting $\tanh(\gamma(c - \bar{z}))$ inside each term, then using Minkowski and Cauchy–Schwarz, gives

$$\int_{\mathbb{R}} |g(t \pm i\vartheta/\gamma)|^2 \, dt \leq \frac{4\gamma}{3} \sec^4 \vartheta \|z - \bar{z}\mathbf{1}\|_2^2 \|u\|_2^2.$$

The derivative bound uses $|\operatorname{sech}^2(t + i\vartheta)| \leq \sec^2 \vartheta \operatorname{sech}^2 t$ and $\int \operatorname{sech}^4 t \, dt = 4/3$ over the real line. Shift the Fourier contour of the analytic function $g(c)^2$ by ϑ/γ; the displayed inequality bounds its absolute integral. Poisson summation and the geometric sum over $|k| > J$ give equation E.12, uniformly for unit v. For a right exterior center, projection cancels the limiting constant vector and $1 - \tanh t \leq 2e^{-2t}$ gives $M^{-1}\|Q^T t_\gamma(c)\|_2^2 \leq 4e^{-4\gamma(c - z_{\max})}$. Summing both tails proves equation E.13. The infinite lattice sum equals $H_\gamma + G_\gamma + R_\gamma$ with $\|R_\gamma\|_2 \leq \delta_\gamma$. Subtracting all absent centers gives $Q^T \Phi_\gamma Q = S_\gamma + R_\gamma - T_{\mathrm{tail},\gamma}$, where $0 \preceq T_{\mathrm{tail},\gamma} \preceq \tau_\gamma I$, proving equation E.14.

E.4 Finite eigenvalue ratios.

Write $\beta_1 \geq \cdots \geq \beta_{M-1}$ for the eigenvalues of the explicit matrix S_γ. They are computed from equation E.5–equation E.11, not fitted to the finite kernel's small eigenvalues. For $e = \mathbf{1}/\sqrt{M}$, define

$$\ell_\gamma = e^T \Phi_\gamma e = 1 + \sum_{j=1}^{W} \left[\frac{1}{M} \sum_{a=1}^{M} \tanh(\gamma(z_a - x_j))\right]^2. \tag{E.15}$$

Then $1 \leq \ell_\gamma \leq \mu_1 \leq W + 1$ by Rayleigh's principle and the trace bound. A sharper upper bound, useful for the two-sided interval, is

$$b_\gamma = \|Q^T \Phi_\gamma e\|_2, \quad a_\gamma = \beta_1 + \delta_\gamma, \quad L_\gamma = \min\left\{W + 1, \frac{\ell_\gamma + a_\gamma + \sqrt{(\ell_\gamma - a_\gamma)^2 + 4b_\gamma^2}}{2}\right\}. \tag{E.16}$$

Indeed, in the basis $[e, Q]$, bound the off-diagonal block by b_γ and the restricted block by $a_\gamma I$ using equation E.14. The resulting scalar 2×2 matrix bounds the largest Rayleigh quotient. Computing $\Phi_\gamma e = A_\gamma(A_\gamma^T e)$ requires no eigendecomposition of the finite kernel.

Theorem E.1 (Finite-network rate bounds)**.** *Under the preceding definitions, set* $\underline{\rho}_1 = \overline{\rho}_1 = 1$*. For* $2 \le i \le M-1$*, define*

$$\underline{\rho}_i(\gamma) = \frac{[\beta_i - \delta_\gamma - \tau_\gamma]_+}{L_\gamma}, \qquad \overline{\rho}_i(\gamma) = \min\left\{1, \frac{\beta_{i-1} + \delta_\gamma}{\ell_\gamma}\right\}, \tag{E.17}$$

where $[s]_+ = \max(s, 0)$*. At* $i = M$*, use* $\underline{\rho}_M = 0$ *and the same upper formula. Then* $\underline{\rho}_i \le \mu_i/\mu_1 \le \overline{\rho}_i$ *for every* i*. For indices* $i > W + 1$*, both endpoints may instead be set to zero.*

Proof. Let κ_i be the ordered eigenvalues of $Q^T \Phi_\gamma Q$. The matrix enclosure and compression interlacing give

$$\beta_i - \delta_\gamma - \tau_\gamma \le \kappa_i \le \beta_i + \delta_\gamma, \qquad \mu_i \ge \kappa_i \ge \mu_{i+1}.$$

Thus the lower bound for μ_i uses β_i, and its upper bound uses β_{i-1}. Divide by the appropriate upper or lower bound on μ_1. Nonnegativity supplies the last lower endpoint. Finally, $\operatorname{rank} A_\gamma \le W + 1$ proves the stated exact zeros. □

The normalization is also explicit in γ. An exterior center at distance $d_j > 0$ from the sample interval contributes at least $\tanh^2(\gamma d_j)$ to equation E.15. These feature means can already be near saturation while the high-frequency factors in equation E.7 remain small. This explains how the numerator and denominator can change at different rates; the finite bounds retain both changes. Increasing γ orders H_γ, but does not by itself order every normalized eigenvalue of Φ_γ. A particular ratio increase is certified whenever $\underline{\rho}_i(\gamma_b) > \overline{\rho}_i(\gamma_a)$. For fixed J, ϑ and exterior-center sets, δ_γ increases and τ_γ decreases with γ; more retained lattice terms can be needed at larger slopes.

E.5 Proof of the output-error bound and conversion to steps.

The upper endpoints $\overline{\rho}_i$ above are denoted ρ_i in the informal statement of Theorem 3.2; here $\rho_i = \mu_i/\mu_1$ denotes the exact ratio. Since $(1-\rho/2)^{2n}$ decreases on $0 \le \rho \le 1$, substituting equation E.17 into equation E.4 proves Theorem 3.2:

$$\underbrace{\sum_i p_i (1 - \overline{\rho}_i/2)^{2n}}_{\underline{E}_\gamma(n)^2} \le E_\gamma(n)^2 \le \underbrace{\sum_i p_i (1 - \underline{\rho}_i/2)^{2n}}_{\overline{E}_\gamma(n)^2}. \tag{E.18}$$

The p_i are the actual finite-kernel target weights from equation E.3, not weights from H_γ or S_γ. For $0 < \varepsilon < 1$, let n_ε be the first integer with $E_\gamma(n) \le \varepsilon$. With $\min \varnothing = +\infty$,

$$\min\{n \ge 0 : \underline{E}_\gamma(n) \le \varepsilon\} \le n_\varepsilon \le \min\{n \ge 0 : \overline{E}_\gamma(n) \le \varepsilon\}. \tag{E.19}$$

Hence upper rate bounds give necessary training times; lower rate bounds give sufficient times. All statements concern the specified frozen-feature, exact-arithmetic gradient descent recurrence.

E.6 Empirical validation of the readout bounds

Figure 3 compares the finite-kernel eigenvalue-ratio bounds with the measured readout learning curves. At fixed center spacing, increasing bandwidth increases $\gamma = \lambda/h$. The left panel checks the ratio intervals from Theorem E.1; the right panel checks the resulting relative-error lower bounds against executed gradient descent.

The error bounds use the target projections $p_i(\gamma)$ of each actual finite kernel, together with the upper ratio bounds $\overline{\rho}_i(\gamma)$. Thus the comparison accounts for changes in both the spectrum and target alignment; no decay rate is fitted to the training trajectory.

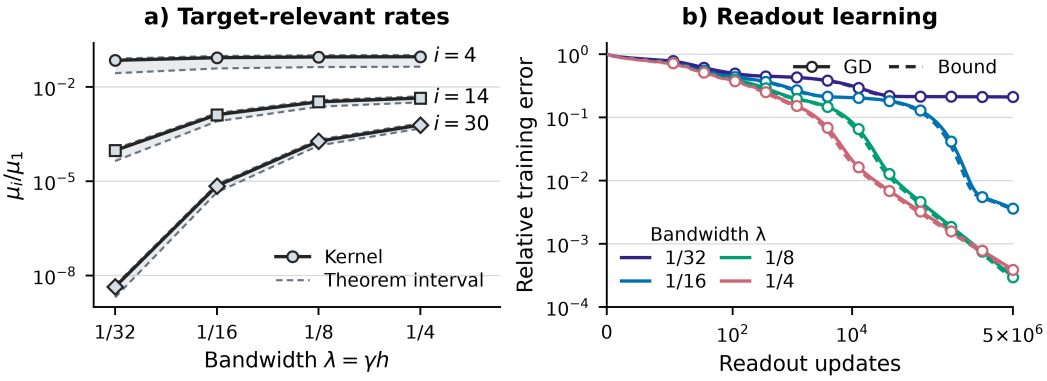


Figure 7: **Finite-kernel rate bounds and readout learning.** Width $W = 512$; target $f(x) = [\sin(2\pi x) + \frac{1}{2}\sin(6\pi x) + \frac{1}{4}\sin(14\pi x)]/\sqrt{21/32}$ on $[-1, 1]$. **Left:** Kernel ratios μ_i/μ_1 (dots) and intervals from Theorem E.1 (dashed boundaries and shading) for three ranks carrying substantial target energy. Lines connect ordered ranks; eigenvectors change with bandwidth. **Right:** Executed frozen-geometry gradient descent errors (solid lines with hollow markers) and relative-error lower bounds (dashed), obtained by taking square roots of the lower bound in equation E.18. The bound uses all numerically resolved target projections onto the frozen kernel's eigenvectors. Errors are relative training RMS through five million updates.

E.7 Numerical example and scope.

The 9.20×10^{11} example uses $N = 128$, $h = 1/64$, 129 core centers and 12 halo centers per side ($W = 153$), 263 equally spaced endpoint-inclusive samples, and $f(x) = \sin(2\pi x) + \frac{1}{2}\sin(6\pi x) + \frac{1}{4}\sin(10\pi x)$. At $\gamma = 4$, the evaluated necessary, finite-kernel spectral, and sufficient counts for 1% relative training error are respectively 9.20×10^{11}, 1.008×10^{12}, and 3.31×10^{12}. A numerical readout refit gives relative error 5.85×10^{-4}. This is a separate configuration from the main width-512 experiment.

For these calculations, a positive quadrature feature factor evaluates H_γ, and retained corrections are added in its singular-vector basis to avoid cancellation in small eigenvalues. Truncating the center integral at distance P/γ beyond the sample interval omits a positive matrix of norm at most $2e^{-4P}/(h\gamma)$; its allowance is added to upper eigenvalue endpoints and the restricted-block bound used in L_γ. Ratios for the actual finite kernel are computed by squaring feature singular values. Powers use $\exp(2n\log(1-\rho/2))$ with `log1p`, and crossing counts use integer bracketing and binary search, not trillion-step training.

The calculation uses $\vartheta = \arctan(\pi/(2\gamma h))$, retains enough lattice terms to make $\delta_\gamma \le 10^{-18}$, and retains $\lceil 18/(\gamma h)\rceil$ absent centers per side. Actual feature-SVD ratios above 10^{-18} are treated as numerically resolved. The saved checks increase quadrature order from 10 to 16 and padding from 20 to 24, and compare independent feature-SVD drivers. For the full error curves, the calculations also use the empirical roundoff allowance $64\varepsilon_{64}\|S_\gamma\|_2$, with $\varepsilon_{64} = 2^{-52}$: subtract it from lower eigenvalue numerators and add it to upper numerators and the restricted-block upper bound. Numerically unresolved components are dropped only from the lower error curve; the upper curve gives them zero lower rate and retains their full energy. These are refinement-checked floating-point evaluations of the analytic bounds, not directed-rounding certificates for the displayed counts.

F SLOW FEATURE ACQUISITION

We formalize the conditional GD statement in Section 3.5 and the Adam acquisition interventions. The structural estimates and the empirical persistence of their premises are separate claims. All output norms below are the training-sample RMS norm $\|\cdot\|_Z := \|\cdot\|_{2,Z}$ on points in $[-1,1]$; parameter norms are Euclidean.

F.1 COUPLED DYNAMICS AND POPULATION BOUNDS

Write $q_\theta(x) = b + \sum_{j=1}^W w_j \tanh(\gamma_j x + \beta_j)$, $e = q_\theta - f$, and $L = \|e\|_Z^2/2$. The hidden biases are free. In gradient flow, for example, $\dot\gamma_j = -w_j\langle e, x\operatorname{sech}^2(\gamma_j x + \beta_j)\rangle_Z$; thus slopes, readouts, and the population residual remain coupled. Let P_C project onto $\operatorname{span}\{1, x\}$, let $P_H = I - P_C$, and write $q_H = P_H q$, $e_C = P_C e$, $e_H = P_H e$, and $f_\perp = P_H f$. Using an orthonormal coarse basis, set $J_C = D(P_C q)$, $J_H = Dq_H$, and $K_C = J_C J_C^*$. When K_C is invertible, define

$$\Pi = I - J_C^* K_C^{-1} J_C, \qquad F = \Pi J_H^* e_H, \qquad R = J_C^*\big(e_C + K_C^{-1} J_C J_H^* e_H\big). \tag{F.1}$$

Here Π is the orthogonal projector onto $\ker J_C$ (Trefethen & Bau, 2022, Lecture 6). Substitution gives $\nabla L = F + R$, $J_C F = 0$, $F \perp R$, and $\|F\| \le \|J_H^* e_H\|$. Importantly, F includes coarse compensation; small R does not discard that response. Full gradient flow dissipates L; effective flow $\dot\theta = -F$ satisfies $\frac{d}{dt}\|e_H\|_Z^2 = -2\|F\|^2$. For native GD, a residual bound is a separate hypothesis, satisfied in particular when the loss never exceeds its restart value.

Population structure and target loading. For $p_j = (\gamma_j, \beta_j, w_j)$, put

$$r_j = \|p_j\|, \qquad M_h = \sum_j r_j^2, \qquad S_k = \sum_j r_j^k, \qquad \chi_k = \frac{W^{k/2-1} S_k}{M_h^{k/2}} \quad (M_h > 0). \tag{F.2}$$

The output bias is excluded. Equal energy on m neurons gives $\sqrt{\chi_6} = W/m$; common expansion changes M_h but not χ_k. For the refined estimates, assume a symmetric training measure on $[-1,1]$ with $\mu_2 > 0$, where $\mu_k = \langle x^k, 1\rangle_Z$. Define $U_2 = x^2 - \mu_2$, $U_3 = x^3 - (\mu_4/\mu_2)x$, $s_i = \|U_i\|_Z^2$, and $\tau_k = \|P_{\mathcal{P}_k} f_\perp\|_Z$, where $\mathcal{P}_k = P_H \operatorname{span}\{1, x, \dots, x^k\}$. Set

$$\begin{gathered}(\alpha_3, \alpha_5, \alpha_7) = (1/3, 2/15, 17/315), \qquad A_p = \alpha_p \frac{2^{p/2} p^{p/2}}{(p+1)^{(p+1)/2}}, \qquad C_p = (p+1)A_p, \\ D_3^2 = \frac{8}{27} s_2 + \frac{3}{16} s_3, \qquad B_3^2 = \frac{s_2}{64} + \frac{3 s_3}{256}.\end{gathered} \tag{F.3}$$

Here $p \in \{3, 5, 7\}$; in particular $A_3 = \sqrt{6}/8$ and $C_3 = \sqrt{6}/2$.

Lemma F.0.1 (Population sensitivity and target loading)**.** *For exact tanh and $\|e_H\|_Z \le E_0$, define $Q = \min\{A_3 S_4,\ B_3 S_4 + A_5 S_6\}$. Then $\|q_H\|_Z \le Q$, and $\|F\|$ is bounded by each of*

$$\begin{aligned} V^{\mathrm{gen}} &= C_3 E_0 \sqrt{S_6}, \\ V^{\mathrm{proj}} &= E_0(D_3\sqrt{S_6} + C_5\sqrt{S_{10}}), \\ V^{\mathrm{tar},3} &= D_3\sqrt{S_6}(\tau_3 + Q) + C_5 E_0\sqrt{S_{10}}, \\ V^{\mathrm{tar},5} &= D_3\sqrt{S_6}\tau_3 + C_5\sqrt{S_{10}}\tau_5 \\ &\qquad + (D_3\sqrt{S_6} + C_5\sqrt{S_{10}})Q + C_7 E_0\sqrt{S_{14}}. \end{aligned} \tag{F.4}$$

The generic force bound and $\|q_H\|_Z \le A_3 S_4$ do not require symmetry.

Proof. The remainder estimates below hold globally, without a Taylor convergence neighborhood. For $u \ge 0$, put $\delta = u - \tanh u$, $v = \tanh u - u + u^3/3$, and $\omega = u - u^3/3 + 2u^5/15 - \tanh u$. Successive use of

$$0 \le \delta' = \tanh^2 u \le u^2, \quad 0 \le v' = \delta(u + \tanh u) \le 2u^4/3, \quad 0 \le \omega' = 2uv + \delta^2 \le 17u^6/45$$

gives $|h_p(u)| \le \alpha_p |u|^p$ and $|h_p'(u)| \le p\alpha_p |u|^{p-1}$ for $h_3 = -\delta$, $h_5 = v$, and $h_7 = -\omega$; parity extends these bounds to all real u. At fixed $r^2 = \gamma^2 + \beta^2 + w^2$, $|\gamma x + \beta| \le \sqrt{2}(\gamma^2 + \beta^2)^{1/2}$.

Maximizing the resulting products in $t = (\gamma^2 + \beta^2)/r^2$ gives output bound $A_p r^{p+1}$ and squared Jacobian-column sum $C_p^2 r^{2p}$ (both maxima occur at $t = p/(p+1)$). Summing output norms and squared column norms therefore gives the generic bounds and the remainder estimates below.

With $u_j = \gamma_j x + \beta_j$, let $q_3 = -\frac{1}{3} P_H \sum_j w_j u_j^3$, $q_5 = \frac{2}{15} P_H \sum_j w_j u_j^5$, and $J_i = Dq_i$. Symmetry makes U_2, U_3 orthogonal, with

$$q_3 = -\Big(\sum_j w_j \gamma_j^2 \beta_j\Big) U_2 - \frac{1}{3}\Big(\sum_j w_j \gamma_j^3\Big) U_3.$$

Differentiation gives

$$\|J_3\|_{\mathrm{HS}}^2 = s_2 \sum_j (4 w_j^2 \gamma_j^2 \beta_j^2 + w_j^2 \gamma_j^4 + \gamma_j^4 \beta_j^2) + s_3 \sum_j (w_j^2 \gamma_j^4 + \gamma_j^6/9).$$

At fixed r_j, the two brackets are at most $8r_j^6/27$ and $3r_j^6/16$; also $w_j^2 \gamma_j^4 \beta_j^2 \le r_j^8/64$ and $w_j^2 \gamma_j^6/9 \le 3r_j^8/256$. Thus $\|J_3\| \le D_3 \sqrt{S_6}$ and $\|q_3\|_Z \le B_3 S_4$. The preceding global estimates give

$$\|q_H - q_3\|_Z \le A_5 S_6, \quad \|J_H - J_3\| \le C_5 \sqrt{S_{10}}, \quad \|J_5\| \le C_5 \sqrt{S_{10}}, \quad \|J_H - J_3 - J_5\| \le C_7 \sqrt{S_{14}}.$$

Finally, split the residual before bounding:

$$J_H^* e_H = J_3^* q_H - J_3^* f_\perp + (J_H - J_3)^* e_H.$$

The columns of J_3 lie in $\mathcal{P}_3$, so its target term is bounded using τ_3, not the full target norm. This gives $V^{\mathrm{tar},3}$ after applying Π. Retaining $J_3 + J_5$ gives $V^{\mathrm{tar},5}$ in the same way. The generic and projected allowances follow directly from the Jacobian bounds. □

In particular, $\|F\| \le C_3 E_0 \sqrt{\chi_6} M_h^{3/2}/W$: at specified total energy, distributed energy limits nonlinear drive. The cubic target-aware allowance is $O(W^{-2})$ when $\tau_3 = 0$ and the moments and energy are bounded, rather than the generic $O(W^{-1})$. These are conditional force scalings, not unconditional acquisition-time bounds.

F.2 Conditional GD comparison

Fix a horizon T, moment allowances $\bar{\chi}_{k,n} \ge \chi_k(\theta_n)$ for $k \in \{4, 6, 10, 14\}$ and $0 \le n \le T$, and tracking allowances $\bar{R}_n \ge \|R_n\|$ for $n < T$. For $B \ge 0$ define

$$\mathcal{S}_{k,n}(B) = \frac{\bar{\chi}_{k,n} B^k}{W^{k/2-1}}, \qquad Q_n(B) = \min\{A_3 \mathcal{S}_{4,n}(B),\ B_3 \mathcal{S}_{4,n}(B) + A_5 \mathcal{S}_{6,n}(B)\}. \tag{F.5}$$

Let $V_n(B)$ be the minimum of the four expressions in (F.4) after substituting $\mathcal{S}_{k,n}(B)$ for S_k and $Q_n(B)$ for Q. Both functions are nonnegative and nondecreasing.

Theorem F.1 (Conditional population-to-output bound for native GD)**.** *Let $f \ne 0$ in the training norm and use the symmetric measure above. Consider $\theta_{n+1} = \theta_n - \eta_n \nabla L(\theta_n)$ with $\eta_n \ge 0$. Suppose $M_h(\theta_n) > 0$, $K_C(\theta_n)$ is invertible, $\|e_{H,n}\|_Z \le E_0$, and the finite structural allowances hold throughout $0 \le n \le T$. Starting at $B_0 = \sqrt{M_h(\theta_0)}$, define*

$$B_{n+1} = B_n + \eta_n\{V_n(B_n) + \bar{R}_n\}. \tag{F.6}$$

Then, for $0 \le n \le T$,

$$\sqrt{M_{h,n}} \le B_n, \qquad \lambda_{\mathrm{RMS},n} = \frac{h\|\gamma_n\|}{\sqrt{W}} \le \frac{hB_n}{\sqrt{W}}, \qquad \frac{\|f - q_n\|_Z}{\|f\|_Z} \ge \frac{[\|f_\perp\|_Z - Q_n(B_n)]_+}{\|f\|_Z}. \tag{F.7}$$

Proof. For $b_n = \sqrt{M_{h,n}}$, the exact GD update and Lemma F.0.1 give $b_{n+1} \le b_n + \eta_n(\|F_n\| + \|R_n\|) \le b_n + \eta_n\{V_n(b_n) + \bar{R}_n\}$. Monotonicity and induction yield $b_n \le B_n$, hence $\|\gamma_n\| \le B_n$ and $\|q_{H,n}\|_Z \le Q_n(B_n)$. Projection and the reverse triangle inequality give $\|f - q_n\|_Z \ge \|f_\perp - q_{H,n}\|_Z \ge [\|f_\perp\|_Z - Q_n(B_n)]_+$. □

Unused force alternatives and their higher-moment hypotheses may be omitted; the generic alternative also applies without symmetry. The proof is a native GD induction, not a long-time GD-to-flow approximation. Endpoint moments enter Q_n; an accumulated sixth-moment condition alone does not supply them. The construction reference $\lambda = 1/4$ is not a necessary accuracy threshold: the output floor is an independent conclusion for the attached readouts.

F.3 Empirical applicability of the GD bound

The reported audit uses full-batch float64 GD, step size 0.002, and 2048 midpoint training points. The primary interval is total updates 25k–125k; updates 20k–25k supply a preceding-window concentration reference. Widths 705 and 1409 include halo units and use $h = 2/512$ and $2/1024$, respectively. On the primary 23-target, two-seed width-705 panel:

Force allowance	Median useful duration	Full 100k intervals
Generic	16.5k	$0/46$
Affine-projected	87.75k	$8/46$
Target-refined	100k	$32/46$

Here "useful" requires both the evaluated upper scale bound below $1/4$ and the training-error floor above 1%. The target-refined comparison lasts at least 79k additional updates in every primary case; projection alone covers all six width-1409 cases through 100k. A factor-two accumulated-$\sqrt{\chi_6}$ diagnostic passes every sampled prefix in all 58 native-GD cases, but the refined evaluations also use measured higher moments and tracking. That diagnostic alone does not verify all their premises. The plotted comparisons use interpolated structural records, not outward-rounded enclosures at every update; halving diagnostic density changes available endpoint bounds by at most 1.17×10^{-4} relatively and durations by at most one 500-update interval. Independent-grid error on 8192 points is a separate empirical quantity. These results are conditional evaluations on the named continuations, not checkpoint-only forecasts or coverage of the separate five-million-update runs.

F.4 Adam: adaptive processing of fine and tracking gradients

The decomposition (F.1) remains exact under Adam, but the GD sensitivity bound is uninformative at its measured higher energies and concentrations. Here we define the actual update accounting and intervention; we do not assert an Adam acquisition-time theorem.

Applied components, not separate optimizers. Use full-batch Adam without weight decay, with the recurrence and bias correction of Kingma & Ba (2015, Algorithm 1). The global update index k and all histories are retained at a fork. For $g_k = F_k + R_k$, linearity of the first-moment recurrence gives $m = m^F + m^R$, where each component uses the same coefficient β_1 and its inherited gradient history. With $D_k = \operatorname{diag}((\sqrt{\widehat{v}_{k+1}} + \epsilon)^{-1})$,

$$u_k^F = -\eta D_k \widehat{m}_{k+1}^F, \qquad u_k^R = -\eta D_k \widehat{m}_{k+1}^R, \qquad \Delta\theta_k = u_k^F + u_k^R, \\ v_{k+1} = \beta_2 v_k + (1-\beta_2)(F_k + R_k)^{\odot 2}. \tag{F.8}$$

Hats use the global bias-correction exponents. The coordinatewise square retains $2F_k \odot R_k$; raw orthogonality does not cancel each coordinate's cross term. Nor does $J_C F = 0$ imply $J_C D_k \widehat{m}_{k+1}^F = 0$. Expanding the squared slope norm gives the exact finite-step identity

$$\|\gamma_{k+1}\|^2 - \|\gamma_k\|^2 = 2\langle \gamma_k, u_{\gamma,k}^F \rangle + 2\langle \gamma_k, u_{\gamma,k}^R \rangle + \|u_{\gamma,k}^F + u_{\gamma,k}^R\|^2. \tag{F.9}$$

The squared-step term must be retained. Define fine activity and coherence by

$$P_F(I) = \sum_{k \in I} \|u_{\gamma,k}^F\|, \qquad \mathcal{C}_F(I) = \frac{\|\sum_{k\in I} u_{\gamma,k}^F\|}{P_F(I)} \quad (P_F(I) > 0).$$

They describe path and net fine displacement; neither alone determines outward growth in (F.9).

Why oscillating tracking can restrict fine amplitude. For illustrative prescribed inputs $F = (\varphi, \varphi)$ and $R_k = (-1)^k(\rho, -\rho)$, the stationary second moment *after processing* g_k is

$$v_{k+1}^{\pm} = \varphi^2 + \rho^2 \;\pm\; 2\varphi\rho \frac{1-\beta_2}{1+\beta_2}(-1)^k. \tag{F.10}$$

Substitution in the EMA recurrence verifies the expression. The fine first moment tends to (φ, φ), while the relative alternating correction in (F.10) is at most $(1-\beta_2)/(1+\beta_2)$. For β_2 near one and negligible ϵ, fine motion per coordinate is therefore approximately $\eta|\varphi|/\sqrt{\varphi^2+\rho^2}$: tracking affects its denominator even without a mean tracking input. This example is not a neural feedback reduction. General cross terms have no fixed sign, so the denominator effect must be tested rather than assumed.

F.5 Adam intervention and plotted quantities

The completed study uses $\eta = 0.002$, $(\beta_1, \beta_2) = (0.9, 0.999)$, $\epsilon = 10^{-8}$, six targets, widths $705, 1409$, and seeds $30, 31$. Alongside the native second moment, evolve

$$v^{\mathrm{sh}}_{k+1} = \beta_2 v^{\mathrm{sh}}_k + (1-\beta_2)(g_k - 0.9R_k)^{\odot 2} = \beta_2 v^{\mathrm{sh}}_k + (1-\beta_2)(F_k + 0.1R_k)^{\odot 2}. \tag{F.11}$$

Initialize the shadow buffer from the stored full buffer at 125k, then run 5k native updates before the common 130k fork. Define D^{sh}_k using the same global bias correction and offset as D_k. For the hidden-parameter selector P_{hid}, let

$$v_N = P_{\mathrm{hid}} u^F_k, \qquad v_S = -\eta P_{\mathrm{hid}} D^{\mathrm{sh}}_k \widehat{m}^F_{k+1}, \qquad \nu = \min\{\|v_S\|, 10\|v_N\|\}.$$

Besides native Adam, compare scalar amplification $v_{\mathrm{gain}} = \nu v_N/\|v_N\|$ and the shadow-denominator proposal $v_{\mathrm{sh}} = \nu v_S/\|v_S\|$. These replace only the hidden fine component; tracking, the total output-bias update, and native moment recurrences retain their native formulas at each clone's current state. The measured proposal norms are nonzero. Local norm matching does not equate cumulative paths or subsequent states. After a 10k pulse, both interventions return to native Adam for a 10k release with inherited moments.

The main-text scatter uses only the 130k–140k pulse. For each policy and matched target–width–seed start, its coordinates are

$$x = \frac{P_F^{\mathrm{intervention}}}{P_F^{\mathrm{native}}}, \qquad y = 100\left(\frac{\|\gamma^{\mathrm{intervention}}_{140\mathrm{k}}\|/\sqrt{W}}{\|\gamma^{\mathrm{native}}_{140\mathrm{k}}\|/\sqrt{W}} - 1\right). \tag{F.12}$$

Thus x is realized fine-slope path, not the nominal gain; y compares the whole-network endpoint with native Adam, not with the fork. The local hidden-update cap does not cap x at ten. There are 48 pulse points when both policies are shown for all 24 starts; release endpoints and error responses are separate diagnostics.

F.6 Adam evidence and limits of the interpretation

Along native paths, tracking attenuation exposes mean capped fine-proposal gains of 5.30–9.95. Scalar amplification increases pulse-end RMS slope in every degree-five, sine, and kink case. However, the width-705, seed-30 compact-bump example accumulates 15.6 times the native fine path with only 0.048% extra endpoint RMS; fine coherence falls from 0.529 to 0.081. The shadow policy gives 73.3 times the path and -1.01% RMS in that case. All altered pulse and release endpoints remain below normalized RMS bandwidth 0.013, compared with the construction reference 0.25; degree-five, sine, and bump errors remain above 1% at those endpoints. Cancellation here can be *induced by amplification*: it is not a universal native-Adam condition.

Coarse-stability diagnostic and discriminating controls. For fixed $D \succ 0$, uncorrected EMA momentum on a positive quadratic mode of $D^{1/2}HD^{1/2}$ is stable when $0 < \eta\lambda < 2(1+\beta_1)/(1-\beta_1)$; we use the standard calculation and preconditioned-sharpness diagnostic of Cohen et al. (2022, Section 2 and Appendix A) rather than repeat its proof. In the separate late-state audit, adaptive sharpness divided by this threshold has median 0.9865 across 24 Adam states; the median coarse contribution to its leading Rayleigh curvature is 0.9880. This uses the full Hessian and actual denominator, but the frozen threshold is not a stability certificate for changing, bias-corrected Adam.

The controls limit stronger causal interpretations. Coarse balancing removes $J_C u^F$ at native amplitude but gives median balanced/native endpoint RMS 1.00012, not a substantial release of acquisition. Turning off active fine motion retains late tracking in degree five, sine, and kink, while some localized cases settle. Keeping coarse output fixed to first order does not keep its compensating equilibrium fixed: with $e^*_C = -K_C^{-1} J_C J_H^* e_H$ and $z = e_C - e^*_C$, the exact identity is $\Delta z = \Delta q_C - \Delta e^*_C$. Finally, freezing the fine denominator at unit gain causes $10/24$ unbalanced and $6/24$ balanced branches to fail; survivor-only comparisons are not all-case effects.

Together, the identities and interventions support a restriction on available fine amplitude and target-dependent responses to increasing it, not one universal source of tracking persistence. Errors use the attached readouts; training-grid versus independent-grid error and current versus best-within-budget error remain distinct. These are empirical results on the named continuations, not an Adam rate theorem or continuous validation of the separate five-million-update experiment.

G Error propagation for compiled arithmetic circuits

Let Φ be an arithmetic circuit computing $f_\Phi : \Omega \to \mathbb{R}$. Suppose affine gates are exact and each multiplication gate is replaced by a neural multiplier with uniform local error at most δ on its prescribed operand domain.

Define $E_v(\delta)$ recursively, with $E_v(\delta) = 0$ at input and constant nodes. For an affine gate $v = c + \sum_i a_i u_i$,

$$E_v(\delta) := \sum_i |a_i| E_{u_i}(\delta).$$

For a multiplication gate $v = uw$,

$$E_v(\delta) := B_w E_u(\delta) + B_u E_w(\delta) + E_u(\delta) E_w(\delta) + \delta,$$

where B_u and B_w are fixed nonnegative bounds on the exact operand magnitudes over Ω, independent of δ.

Fix $\delta_{\max} > 0$ such that all exact and approximate intermediate values remain in their prescribed domains for $0 < \delta \le \delta_{\max}$. Let o be the output node and define

$$\kappa_\Phi := \frac{E_o(\delta_{\max})}{\delta_{\max}}.$$

Then

$$\|N_{\Phi,\delta} - f_\Phi\|_{\infty,\Omega} \le \kappa_\Phi \delta, \qquad 0 < \delta \le \delta_{\max}.$$

Proof. Induction over a topological ordering of the circuit shows that $E_v(\delta)$ bounds the error at every node. Each $E_v(\delta)$ is a polynomial in δ with nonnegative coefficients and zero constant term, so $E_o(\delta)/\delta$ is nondecreasing. Hence

$$\frac{E_o(\delta)}{\delta} \le \frac{E_o(\delta_{\max})}{\delta_{\max}} = \kappa_\Phi.$$

The size and depth bounds follow by replacing each multiplication gate with the chosen neural multiplier; independent gates at the same circuit level are evaluated in parallel.

G.1 Boolean-circuit compilation

We give the construction underlying Section 4.3. The result concerns representation and execution of a supplied Boolean circuit; the network weights are constructed explicitly rather than learned.

Theorem G.1 (Boolean-circuit compilation)**.** *Let $C : \{0,1\}^n \to \{0,1\}$ be a depth-d Boolean circuit with AND and OR gates, maximum fan-in F, and negations only on input literals. For any $0 < \eta \le 1/4$, there exists an explicitly constructed feed-forward tanh network N_C such that*

$$\sup_{x \in \{0,1\}^n} |N_C(x) - C(x)| \le \eta.$$

Consequently,

$$\mathbf{1}\{N_C(x) \ge 1/2\} = C(x) \qquad \text{for every } x \in \{0,1\}^n.$$

The nonlinear depth is d. In particular, any polynomial-size, constant-depth AC^0 circuit family admits a polynomial-size, constant-depth tanh-network realization with exact correctness after Boolean decoding.

Proof. Set

$$\epsilon = \min\left\{\eta, \frac{1}{8F}\right\},$$

and define the scalar threshold unit

$$H_a(t) = \frac{1 + \tanh(at)}{2}.$$

Choose

$$a = \frac{4}{3} \log \frac{1}{\epsilon}.$$

For $t \geq 3/8$,

$$1 - H_a(t) = \frac{1}{1 + e^{2at}} \leq e^{-2at} \leq e^{-3a/4} = \epsilon,$$

and by symmetry $H_a(t) \leq \epsilon$ whenever $t \leq -3/8$.

Consider an m-input gate, with $m \leq F$, whose intended Boolean inputs are $b_i \in \{0, 1\}$. Suppose the corresponding network activations satisfy

$$|y_i - b_i| \leq \epsilon \qquad \text{for all } i.$$

Then

$$\left| \sum_{i=1}^{m} y_i - \sum_{i=1}^{m} b_i \right| \leq m\epsilon \leq F\epsilon \leq \frac{1}{8}.$$

For an OR gate, use

$$\widehat{\mathrm{OR}}_m(y) = H_a\left(\sum_{i=1}^{m} y_i - \frac{1}{2} \right).$$

On exact Boolean inputs, the argument of H_a is at most $-1/2$ when the correct output is 0 and at least $1/2$ when the correct output is 1. After perturbation by at most $1/8$, it therefore retains the correct sign and has magnitude at least $3/8$. Hence

$$\left| \widehat{\mathrm{OR}}_m(y) - \mathrm{OR}_m(b) \right| \leq \epsilon.$$

Similarly, an AND gate is realized as

$$\widehat{\mathrm{AND}}_m(y) = H_a\left(\sum_{i=1}^{m} y_i - m + \frac{1}{2} \right).$$

Its exact threshold argument is $1/2$ when all inputs are 1 and at most $-1/2$ otherwise. The same perturbation argument therefore gives

$$\left| \widehat{\mathrm{AND}}_m(y) - \mathrm{AND}_m(b) \right| \leq \epsilon.$$

Negated input literals are realized exactly by the linear map

$$\neg x_i = 1 - x_i.$$

Thus every AND or OR gate maps inputs within ϵ of their intended Boolean states to an output within ϵ of the correct Boolean value. Since the input literals are exact, induction over the circuit layers gives

$$|N_C(x) - C(x)| \leq \epsilon \leq \eta$$

for every Boolean input x.

Because $\eta \leq 1/4 < 1/2$, thresholding the final activation at $1/2$ recovers the exact Boolean output:

$$\mathbf{1}\{N_C(x) \geq 1/2\} = C(x).$$

Each source circuit layer contributes one nonlinear tanh layer, so the nonlinear depth is d. Each Boolean gate is represented by one tanh threshold unit together with affine wiring. Hence a circuit family of polynomial size and constant depth yields a tanh-network family of polynomial size and constant depth. □

Weight magnitudes. The required slope satisfies

$$a = O\left(\log \frac{F}{\eta} \right).$$

The largest affine coefficient arises from the bias of an m-input AND gate, whose magnitude is $O(am)$; hence the maximum weight magnitude is

$$O\left(F \log \frac{F}{\eta} \right).$$

Scope. This construction uses tanh threshold units directly rather than the QUILL approximation theorem. It establishes representation and execution of a supplied Boolean circuit. The sorting and shortest-path experiments in Section 4.3 are finite compiled programs and are not asymptotic AC^0 claims.