Supplementary Materials

Neural Networks for Finite-Rate Chemistry in Combustion Simulations: Physical Requirements, Predictive Capability, and Cost

Meng Wang and Tianfang Xie

Numbers in square brackets refer to the reference list of the paper; equations, sections, figures, and tables without the prefix S are those of the paper. Tables S1 to S3 are given in the workbook Supplementary_Tables_S1-S3.xlsx.

S1 Representative coupled flame simulations

Table S4 describes representative two- and three-dimensional flame simulations in which a network computed the chemistry, grouped by the physical problem they test; all 63 such studies are listed in Table S1.

Table S4: Representative coupled flame simulations with neural-network chemistry, grouped by the physical problem they test. Fuel: fuel and reference mechanism (number of species). Network: what the network gives the solver and its training data. DI: the same simulation with direct integration of the chemistry.

Configuration and flow model	Fuel	Network	Main result	Ref.
Ignition and premixed propagation				
Lean H_2/air kernel ignition and flame propagation in decaying turbulence; 2D DNS	H_2; 9 species	state map; stochastic mixing reactors	stand-alone network missed the heat-release peak of ignition and its heat release showed spikes during propagation; a hybrid returning states to CVODE when the sum of the mass fractions deviated from one by more than 2×10^{-4} removed both	[88]
Premixed DME/air flame at 5 atm, 1D; penta-brachial DME flame, 2D	DME; skeletal (39)	state map; unconstrained, penalized, or with exact element conservation	unconstrained map diverged after 0.16 ms, penalized map after about 6 ms; with exact conservation stable over 10 ms, flame speed within 4.9%	[75]
Premixed C_2H_4 and Jet-A flame kernels in decaying turbulence; 2D DNS	C_2H_4 (24); Jet-A (41)	state map; flamelets, augmented	turbulent burning velocity of the C_2H_4 kernel within 0.72% of DI	[39]
Cambridge stratified flames without swirl; flame-resolved LES, 12 million cells	CH_4; DRM19 (20)	state map; flamelets and random states	trend of the rise of the equivalence ratio near the bluff body from differential diffusion captured but much smaller than measured; minor species within 5% of CVODE; displacement-speed statistics of the burnout zone missed	[26]
Extinction and re-ignition				
CH_4/air counterflow flame strained to extinction, 2D; planar jet with local extinction and re-ignition, 3D; DNS	CH_4; GRI-Mech 3.0 (53)	state map; randomized flame states; 0, 1, 5, or 10 unrolled steps	networks trained with ten unrolled steps captured the extinction and the NO of the counterflow flame; one such network captured the conditional statistics of the jet; without unrolling the peak conditional mean temperature of the jet was about 50 to 150 K low	[94]
Sandia flames D, E, F; LES-PDF	CH_4; GRI-Mech 1.2 without Ar (31)	state map; random states generated around flamelet states	temperature and major species nearly identical to DI; local extinction and re-ignition of flame F seen in snapshots	[73]
Mixing and turbulent flame structure				
Premixed NH_3/air jet in burned gas; 2D DNS	NH_3; Okafor et al. (59)	neural ODE; homogeneous ignition	departs from DI after 0.15 ms as burned and unburned gas mix	[106]
Premixed NH_3/air and NH_3/H_2/air flames, Karlovitz number 17 to 437; 2D DNS	NH_3 and NH_3/H_2; 32 species	state map; flamelets, augmented	temperature error below 80 K at moderate Karlovitz number, pockets up to 750 K at high Karlovitz number where the flame position deviated, mean and conditional profiles in agreement	[112]
Heat loss and walls				
Syngas/O_2 jet flame along a cooled wall; 2D and 3D DNS	syngas; GRI-Mech 3.0 without N chemistry (35)	state map (convolutional); target simulation	2D: mean wall heat flux within 3% of GRI-Mech 3.0, standard deviation 6.5% of the flux; 3D: 14% from the reduced scheme used as reference	[139]
PRECCINSTA swirl combustor with wall heat loss and radiation; LES-PDF	CH_4; GRI-Mech 1.2 without Ar (31)	state map; burner-stabilized and prescribed-temperature flames with heat loss	mean temperature and major species very close to DI; HO_2 and C_2H_6 about 10% off	[27]
Elevated pressure and real fluids				
Lean n-heptane/air autoignition with temperature fluctuations at 40 atm; 2D DNS with principal-component transport	n-C_7H_{16}; reduced (58)	principal-component source terms; target simulation	cool flame, ignition kernels, and heat-release statistics reproduced; an initial temperature 50 K above training gave a noticeable ignition error	[125]

Table S4 (continued)

Configuration and flow model	Fuel	Network	Main result	Ref.
Supercritical H_2/O_2 coaxial jet at 7 MPa; LES	H_2 (10); CH_4 (17)	state maps for chemistry and real-fluid properties	mean radial profiles match DI; flame-length error of a supercritical CH_4/O_2 kernel below 5%	[108]
Pollutants				
Sydney swirl flame SMA2 and Sandia flame D; LES-PDF	CH_4; GRI-Mech 3.0 without Ar (52)	state map; flamelets, augmented	first network set: errors in nitrogen intermediates grew in the recirculation zone and corrupted NO, N_2O, and NO_2 (Section 6.2 of the paper); a second set with networks for small concentrations matched DI except for overpredicted peaks of several species near the bluff body	[123]
Pilot-scale H_2 furnace; LES	H_2; 38 species	NO source term added to 9-species chemistry	NO emission index 40% higher with reduced chemistry than with the detailed mechanism, whose value lay inside the measured range; one-way coupling (preprint)	[90]

S2 State of practice

This section collects the counts behind the statements of the paper. All refer to the 125 studies of Appendix A and are computed from the codes, not from free text.

S2.1 Growth and forms

Of the 125 studies, 28 appeared before 2020 and 97 between 2020 and September 2026 (Figure S1). By what the network gives the solver, 76 (61%) use state maps, 36 (29%) source terms, 5 trajectories, and 7 auxiliary decisions, and one study solves the kinetic equations directly with a physics-informed network, counted with the trajectories in Table S5. The one operator network run in a flow solver [89] is counted with the state maps, because it was used as one. Of the 76 state maps, 61 are trained for one or a few fixed time steps.

Table S5: What the network gives the solver in the 125 studies, with the variables in which it works (full state/subset of species/latent coordinates), the treatment of the time step, and the number of studies that ran the network inside a flow solver.

Role	What the network returns	Studies	Variables	Time step	Coupled	Examples
State map	state, or its increment, after a time step	76	62/12/2	fixed 61; input 15	45	[11, 69, 73, 74, 96, 75]
Source terms	instantaneous source terms, integrated by a conventional solver	36	13/8/15	integrator	17	[33, 71, 76, 103, 77]
Trajectory	solution over a time window from the initial state, including one direct physics-informed solution of the kinetic equations	6	5/0/1	window	0	[78, 79, 80, 81]
Auxiliary decision	choice of solver, mechanism, species, or tolerance; states for the reference solver	7	3/2/2	n/a 6; input 1	4	[82, 83, 86, 85]

S2.2 Levels of testing

The tests in the studies follow a common sequence in which each level adds a physical process that can make a network fail and that the previous level cannot reveal: single steps on held-out states, which measure interpolation; recursive 0D runs, in which the network is fed its own output through ignition and relaxation to equilibrium; coupled laminar flames, in which diffusion brings in states that no homogeneous reactor visits; and coupled turbulent simulations, in which mixing of burned and unburned gas, strain, and heat loss widen the composition space further. Of the studies, 10 tested single steps only, and for 49 recursive 0D runs were the most demanding test. The other 66 ran the network inside a flow solver: 2 in laminar flames, 59 in turbulent simulations, and 5 in a practical combustor, one of them only for 100 time steps to measure performance [14] (Figure S2). Coupled tests were the rule in the early work (16 of the 24 studies up to 2015), became less common in 2016 to 2022 (15 of 34), and are used in about half of the recent studies (35 of 67 since 2023). Table S6 compares the studies before and from 2023.

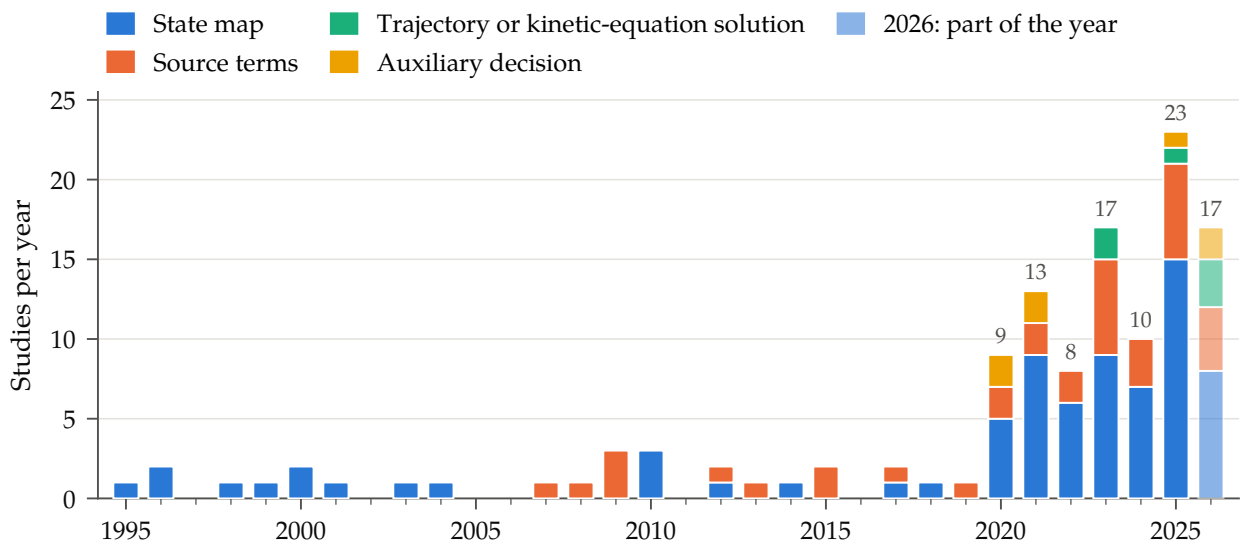


Figure S1: Studies per year, by what the network gives the solver; from 2020, the number above a bar is the total of the year. The search was run in September 2026, so the lighter bar for 2026 covers only part of that year.

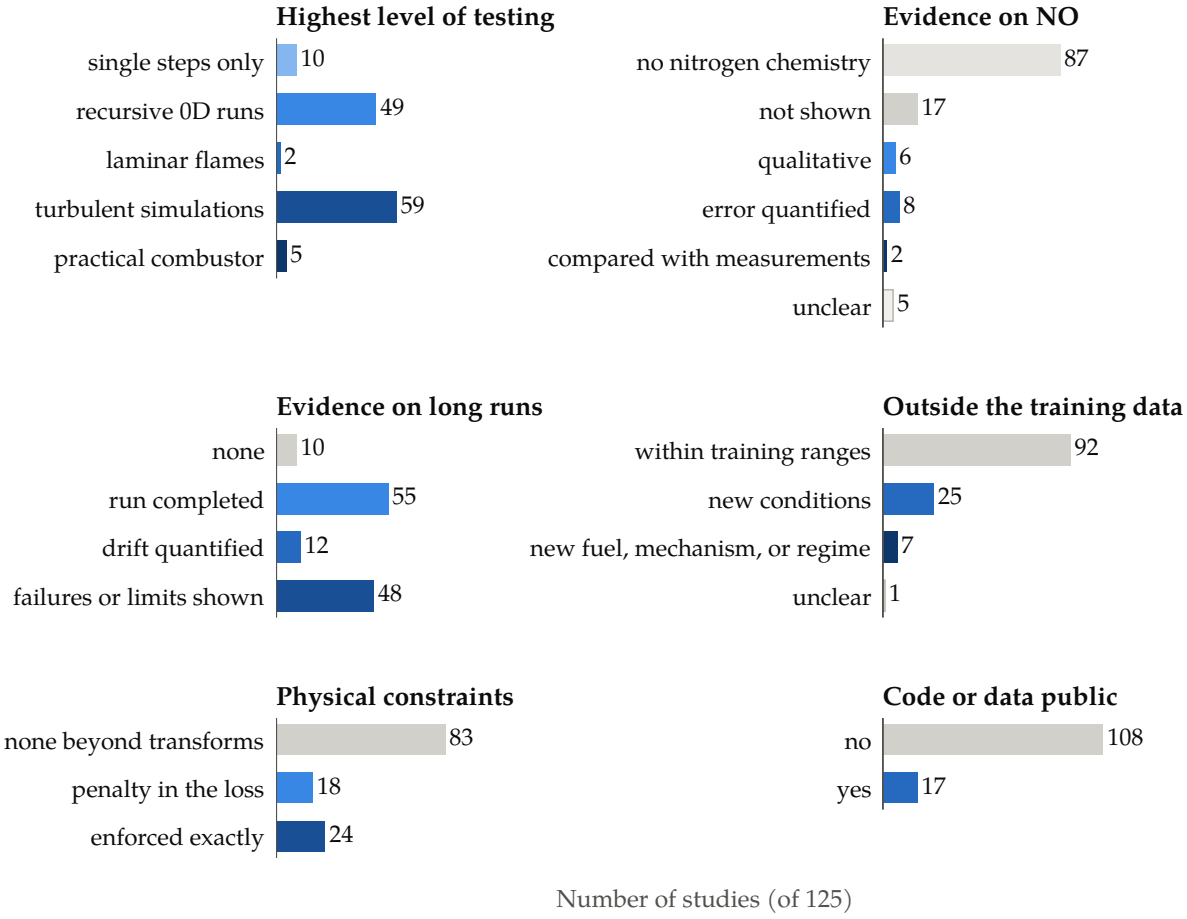


Figure S2: Evaluation practice of the 125 studies, coded from the full texts: number of studies at each level of six criteria. Gray: no evidence on the criterion (lightest gray: mechanism without nitrogen chemistry, for which the NO criterion does not apply); blue, from light to dark: increasingly demanding evidence; unclear: the full text did not allow a decision.

Table S6: Share of studies meeting selected criteria before 2023 (58 studies) and from 2023 (67 studies).

Criterion	Before 2023	From 2023
Network run inside a flow solver	53%	52%
Error growth over a run quantified or failures shown	38%	57%
Conservation enforced by the loss or exactly	21%	45%
Tested outside the training ranges	16%	34%
NO error quantified, studies with nitrogen chemistry	17%	38%
Code, models, or data public	5%	21%

S2.3 Training data, test data, and constraints

The training states come from five kinds of sources, which some studies combine: 57 use homogeneous reactors, 32 laminar flames, 26 random or perturbed states, 13 stochastic mixing reactors, and 24 states from the simulation on which the network is later tested or from a precursor of it; 22 add synthetic states. The size of the training set is stated in 87 studies: the median is 800,000 states, half of the sets lie between about 76,000 and 3.1 million states, and the extremes are 1,000 and 50 million. Of the studies, 37 drew their test states at random from the same pool as the training states, 42 held out whole trajectories, initial conditions, or parameter values, and 29 tested on another configuration; 10 used no separate test set and 7 do not say how it was formed. Not counting the computation of the temperature from the enthalpy, 83 studies use only transforms and normalization, 18 add conservation terms to the loss without any exact constraint, and 24 enforce at least one constraint exactly. By the property constrained, 20 studies penalize the element balance in the loss and 9 enforce it exactly, 14 make the mass fractions sum to one, 10 enforce positivity beyond a transform, and 17 impose energy conservation, including the computation of temperature from the enthalpy. Only 18 studies report the conservation error over a recursive or coupled run, in numbers or curves.

S2.4 Long runs, transfer, combustion targets, and slow species

Of the studies, 55 state only that their runs were completed, 12 show how the error evolves over the run, and 48 report failures or limits of their networks; 17 keep the reference solver in the loop by a rule, an error indicator, or design. Of the studies, 92 tested their networks only within the ranges of their training data, 25 under conditions outside those used to generate the training data, which can still lie inside randomized or perturbed training ranges, and 7 with a new fuel, mechanism, or combustion regime; one study could not be coded. Ignition delays were compared after recursive or coupled use in 59 studies, 14 of them with a stated error; laminar flame speed, flame structure, or turbulent flame propagation in 36 (16); extinction or re-ignition in 10 (4), of which one computed laminar extinction and the others compared turbulent flames with local extinction; and heat release, temperature fields, or flame statistics in turbulent or multidimensional simulations, including stochastic mixing reactors and one-dimensional turbulence models, in 68 (33). In 22 studies, the results were also compared with measurements. Mechanisms without nitrogen chemistry are used in 87 studies, and for 5 the mechanism is not described well enough to tell. Of the 33 studies whose mechanism contains it, 17 do not show NO, 6 show it qualitatively, 8 quantify its error, in single steps or after recursive or coupled use, and 2 compare it with measurements; 11 show NO after recursive or coupled use, 6 of them with a quantified error. CO is quantified in 25 studies.

S2.5 Physical conditions of the tests

Low-temperature chemistry was examined in 12 studies and differential diffusion in 10. Non-adiabatic states were part of the training data or the tests in 13 studies and part of the tests in 8, and wall heat flux or near-wall quenching was compared after coupled use in 2, 1 of them with a stated error. Tests at 10 bar or more appear in 22 studies and real-fluid thermodynamics in 2. Engine conditions are addressed by 5 studies through constant-volume DNS, RANS of a spray chamber, or a priori tests; none of them simulates an engine. Figure S3 shows the distribution of studies that compared each combustion quantity after recursive or coupled use, by fuel.

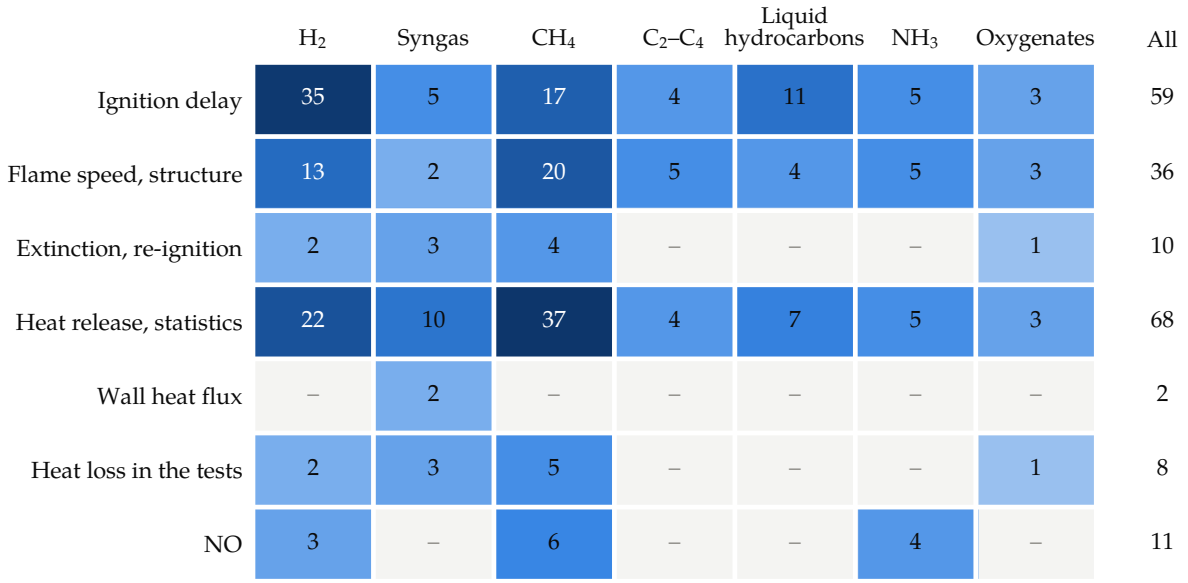


Figure S3: Number of studies that compared each quantity after recursive or coupled use, by fuel class; for heat loss, studies whose tests included non-adiabatic states, and for NO, studies that show NO after recursive or coupled use. A study with several fuels is counted under each; All: number of distinct studies, including those with global or other chemistry. Darker cells: more studies.

S2.6 Reporting and evaluation items

Table S7 lists all reporting and evaluation items, the physics each tests, and the claim that makes it necessary; Table S8 shows, for the items that the codes capture, the criterion coded and the number of studies that meet it. Because the coding does not capture every item in full, a study counted for an item may still miss part of what the item asks.

Table S7: Reporting and evaluation items for neural-network chemistry, the physical reason for each, and when it is needed; the five priorities of Section 8.3 are drawn from this list.

Item	What to report	Physical reason	Needed when
Reference solution	Mechanism (species, reactions), integrator, tolerances, and time step used to generate the data and to judge the network	the error is defined against it	always
Training states	Sampling method; ranges of temperature, pressure, equivalence ratio, and, where relevant, strain, dilution, and enthalpy defect; number of states; test data separated by condition	accessible composition space	always
Single-step errors	Errors per species, normalized by the state and, for the radicals of the induction and the slow species, by the true increment of the step, with a floor near equilibrium where the increment vanishes; relative errors of the species that carry the explosive mode	slow and explosive modes respond to the error relative to the increment	state maps; ignition or slow species claimed
Transfer	At least one test outside the sampled ranges, with the distance from them	coverage of the target flow	transfer is claimed
Conservation	Treatment of element and enthalpy (or energy) conservation, positivity, and the sum of mass fractions, and the remaining element and enthalpy errors after a long run	conserved directions are not damped	always
Coupled tests	The network inside a flow solver, in configurations of increasing dimension, against the same simulation with the reference solver	transport creates new states	use in CFD is claimed
Long runs	Run length in steps, error against the number of calls along a fluid element, failure cases and their cause	recursive error growth	use in CFD is claimed
Combustion targets	Errors in the combustion quantities the network is meant to predict, among ignition delay, flame speed, extinction, and heat release, inside and outside the training ranges	explosive modes and reaction-diffusion balance	for the quantities the study claims to predict
Heat loss and walls	Errors in states with enthalpy defect and, where relevant, in the wall heat flux	enthalpy dimension of the state	walls or radiation matter

Table S7: (*continued*)

Item	What to report	Physical reason	Needed when
Slow species	NO, CO, and other pollutants against the reference, with the spread between mechanisms and measurements where available	slow modes; small concentrations	emissions are claimed
Safeguards	The rule or error indicator that returns states to the reference solver and the share of states it returns	errors amplified at ignition	production use
Fair reference	At least one comparison with an adaptive stiff integrator at a stated tolerance, on equal hardware and with the same mechanism; a GPU integrator when the network runs on a GPU	separates the gain of the method from those of the hardware and the mechanism	a speedup is claimed
Cost	Speedup of the reaction step and of the simulation, chemistry share of the run time with both methods, hardware, batch size, and data transfer; cost of data and training	Equations (9) and (10)	a speedup is claimed
Availability	Code, trained weights, normalization constants, and the data or the scripts that regenerate them	the approximation can be reproduced and tested	always

Table S8: Coded criteria for the items of Table S7 and the number of studies that meet them. Speedup items refer to the 92 studies that report a speedup; the NO item refers to the 33 studies whose mechanism contains nitrogen chemistry.

Item	Coded criterion	Studies
Training states	test data held out by condition or configuration	71 of 125
Transfer	tested outside the training ranges	32 of 125
Conservation	conservation error over a run given in numbers or curves	18 of 125
Coupled tests	network run in a flow solver	66 of 125
Long runs	error growth quantified or failures shown	60 of 125
Combustion targets	at least one target quantified a posteriori	55 of 125
Slow species	NO quantified; CO quantified	NO: 10 of 33; CO: 25 of 125
Safeguards	reference solver in the loop	17 of 125
Fair reference	at least one like-for-like value	39 of 92
Cost	hardware stated for every value; chemistry and total speedup both given	67 of 92; 29 of 92
Availability	code, models, or data public	17 of 125

S3 Reported speedups and mechanism size

Pooled over the studies, 234 speedup values from 92 studies range from a slowdown to about 10^5 (Figure S4). The median chemistry speedup is 10 when the network is compared with an adaptive stiff integrator using the same mechanism (56 values from 35 studies; 10 for the 47 values with stated, identical hardware), but 186 when a network on a GPU is compared with a reference on CPU cores (22 values from 14 studies). These medians pool different fuels, mechanisms, and configurations.

Across the 35 studies that give a like-for-like chemistry speedup and the number of species the network represents (Figure S5), the rank correlation between the two is 0.52 (Spearman coefficient; one-sided permutation test, $p < 0.001$), and the median speedup rises from 3 for mechanisms of up to 13 species (13 studies) to 7.6 for 14 to 30 species, 15 for 31 to 60, and 35 for more than 60 (4 studies). For the smaller mechanisms, the study medians span more than three orders of magnitude, from 1.1 to 2,750 for up to 13 species and from 1.65 to 10,000 for 14 to 30 species; for the larger ones, they span one to two orders of magnitude. This describes the values collected, not a scaling law. Of the 120 studies that state the number of species, 57% use at most 20 species and 8% more than 100.

S4 Coverage of previous reviews

Table S9 codes the thirteen reviews closest to this topic with the same questions as the primary studies (Appendix A.5 of the paper).

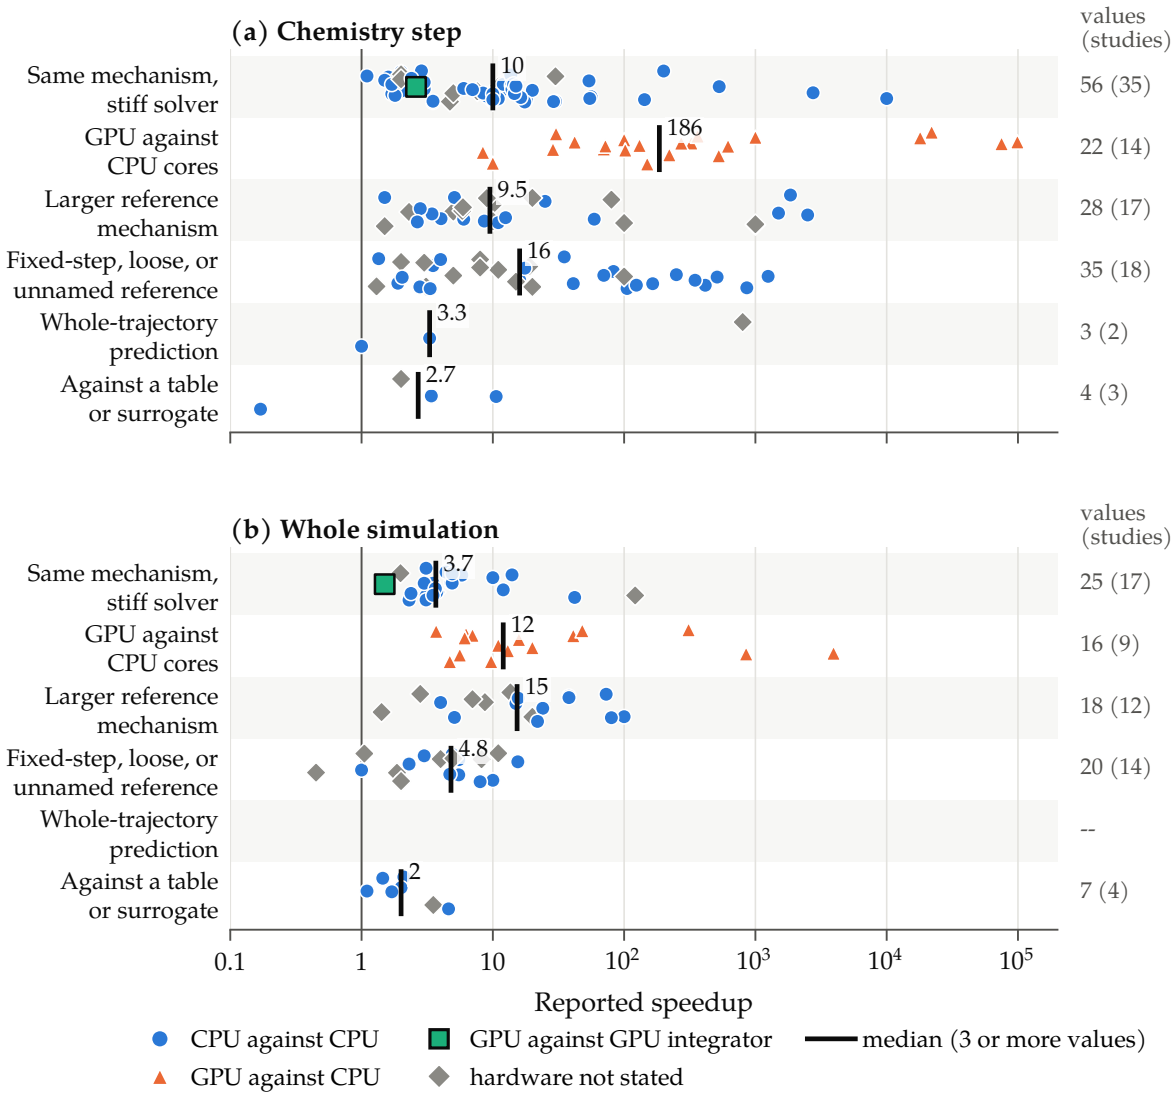


Figure S4: Reported speedups of the 92 studies that give one, for (**a**) the reaction step and (**b**) the whole simulation, grouped by the comparison behind each value. First row, gray markers: hardware not stated; outlined squares: a network against a GPU integrator on the same hardware; black ticks: medians of groups with at least three values; numbers at the right: values (studies) per group, a study with several comparisons counting in several groups; vertical line: equal cost.

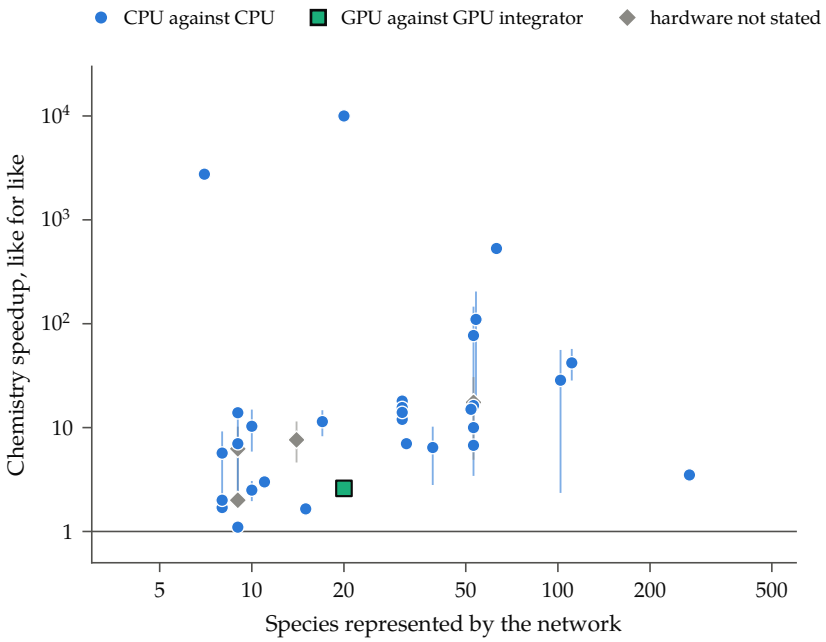


Figure S5: Like-for-like chemistry speedup against the number of species represented by the network, one point per study (median of its values; vertical lines: range of its values).

Table S9: Coverage of neural-network chemistry in previous reviews, coded from their full texts with the same questions as the primary studies. Search: literature search described; Studies: number of primary studies of neural-network chemistry cited; Constr.: physical constraints; Long runs: error accumulation over long runs; Slow species: NO, CO, and other slow or minor species; Outside data: tests outside the training data; Cost basis: the comparison behind reported speedups, hardware, and training cost; Protocol: reporting standard proposed; Grading: evidence of the cited studies tabulated or graded. ●: discussed in at least one paragraph specific to neural-network chemistry; ○: one or two sentences; –: not discussed. [a] Preprint.

Review	Search	Studies	Constr.	Long runs	Slow species	Outside data	Cost basis	Protocol	Grading
Turányi and Tomlin 2013 [162]	–	10	–	–	○	○	○	–	–
Zhou et al. 2022 [15]	–	22	○	–	○	○	○	–	○
Ihme et al. 2022 [16]	–	17	●	–	○	●	○	○	–
Echekki et al. 2023 [17]	–	22	○	–	○	○	○	–	–
Ihme and Chung 2024 [159]	–	7	–	○	–	–	○	–	–
Shejan et al. 2024 [153]	–	15	–	–	–	●	○	–	–
Deng et al. 2025 [158]	–	16	●	–	○	●	○	○	–
Tao et al. 2025 [161]	–	15	●	–	○	○	○	–	–
Wu et al. 2025[a] [163]	–	38	●	–	–	○	○	○	○
Quaye et al. 2025 [160]	–	3	○	–	–	○	○	–	–
Zhao et al. 2026 [164]	–	15	○	○	○	–	○	–	○
Mirgolbabaei 2026[a] [19]	●	14	●	●	○	●	●	●	●
Shateri et al. 2026[a] [18]	–	32	●	○	○	●	●	●	○

S5 Accuracy required of the reaction step: all conditions

Table S10 extends Table 1 of the paper to all conditions computed (methods in Appendix A.3); the data, including the advance for every error, are available from the authors with the scripts (Data Availability Statement of the paper). With the H/O/N subset of the mechanism of Glarborg et al. [1], at 1000 K and 1 to 20 bar, the hydrogen ignition advance for a bias of 10^{-12} per step was 1.3 to 1.5 times that with the mechanism of Stagni et al. [37], and the change of the delay for a 5% relative error was the same within 0.1 percentage points; for ammonia/hydrogen, the advance for the same bias was 3 to 15 times larger. Figure S6 shows the methane ignition delays and their advance for a fixed bias over a wider range of initial temperature.

The $\lambda_e \tau$ fitted to the computed advances with Equation (4) is an effective parameter. For methane, and for hydrogen below the extended second explosion limit with inductions of 40 steps or more, it lies within 28% of the integral of the explosive eigenvalue from 0 to 0.9τ, about 5 for methane and 13 to 17 for hydrogen, consistent with a bias that grows with the explosive mode. For ammonia/hydrogen it is 1.4 to 4.4 times this integral, and for hydrogen above the limit 0.8 to 2.5 times, where the explosive mode is not carried by the species to which the bias is added.

Where the hydrogen induction lasts only a dozen steps of 1 µs, the tolerances also depend on the step (Table S11). At 1200 K and 5 and 10 bar, the bias tolerance per microsecond is 1.8 to 2.3 times smaller at 0.1 µs than at 1 µs, and the delay change for a 5% relative error rises from 2.5% to 3.9% at 10 bar, toward the 4.8% of Equation (6); at 1000 K and 1 bar, with 201 steps, the tolerance is only 3% smaller and the change rises only from 4.2% to 4.4%.

Table S10: Accuracy required of the reaction step at Δt = 1 µs for all conditions computed (ϕ = 0.5 in air); columns as in Table 1 of the paper, with $\lambda_e \tau$ fitted to the computed advances and the integral of the explosive eigenvalue from 0 to 0.9τ; NO bias for 10%: error per step in the NO mass fraction that gives an NO error of 10% at the end of a 10 ms interval starting at 1.5τ.

Mixture	T_0 **(K),** p **(bar)**	τ **(ms)**	$\tau/\Delta t$	$\lambda_e \tau$**, fit**	$\int \lambda_e \, \mathrm{d}t$	**Bias for 5%**	**Delay change for 5% relative error (%)**	**NO bias for 10%**
CH_4/air	1100, 5	28.3	28322	6	5	8.1×10^{-12}	<0.01	7.6×10^{-9}
	1100, 10	13.7	13721	6	5	1.5×10^{-11}	<0.01	8.0×10^{-9}
	1100, 20	6.89	6894	6	5	3.0×10^{-11}	<0.01	9.7×10^{-9}
	1200, 5	7.17	7168	6	5	6.3×10^{-11}	<0.01	1.4×10^{-8}
	1200, 10	3.55	3548	5	5	1.1×10^{-10}	<0.01	1.8×10^{-8}
	1200, 20	1.77	1767	5	5	2.3×10^{-10}	<0.01	2.5×10^{-8}
H_2/air	1000, 1	0.201	201	18	17	7.6×10^{-13}	4.2	4.0×10^{-9}
	1000, 2	0.225	225	13	14	8.8×10^{-13}	4.1	6.1×10^{-9}
	1000, 5	7.90	7899	10	8	3.1×10^{-12}	0.032	1.4×10^{-8}
	1000, 10	8.13	8125	10	8	8.6×10^{-12}	<0.01	2.2×10^{-8}
	1000, 20	5.92	5920	8	7	2.4×10^{-11}	<0.01	3.1×10^{-8}
	1100, 1	0.0857	86	19	17	1.1×10^{-11}	4.0	1.0×10^{-8}
	1100, 2	0.0485	49	17	16	2.3×10^{-11}	3.7	1.6×10^{-8}
	1100, 5	0.0460	46	13	13	3.0×10^{-11}	3.6	2.9×10^{-8}
	1100, 10	0.566	566	9	7	4.6×10^{-11}	0.14	4.6×10^{-8}
	1100, 20	0.692	692	18	7	1.5×10^{-10}	0.064	7.1×10^{-8}
	1200, 1	0.0503	50	19	16	1.1×10^{-10}	3.8	2.4×10^{-8}
	1200, 2	0.0250	25	17	16	2.6×10^{-10}	3.1	3.9×10^{-8}
	1200, 5	0.0119	12	20	14	6.7×10^{-10}	2.5	7.3×10^{-8}
	1200, 10	0.0107	11	22	12	1.1×10^{-9}	2.5	1.2×10^{-7}
	1200, 20	0.0712	71	6	7	5.4×10^{-10}	0.42	1.8×10^{-7}
NH_3/H_2/air	1000, 1	5.05	5049	44	13	1.2×10^{-11}	0.29	1.4×10^{-7}
	1000, 2	11.3	11262	23	10	3.6×10^{-11}	0.11	1.2×10^{-7}
	1000, 5	22.0	22019	13	7	1.0×10^{-11}	<0.01	8.9×10^{-8}
	1000, 10	16.7	16656	11	7	2.2×10^{-11}	<0.01	7.4×10^{-8}
	1000, 20	10.2	10219	10	7	4.2×10^{-11}	<0.01	5.9×10^{-8}
	1100, 1	1.19	1187	36	12	4.2×10^{-11}	0.55	1.7×10^{-7}
	1100, 2	1.07	1066	43	11	1.6×10^{-10}	0.33	1.5×10^{-7}
	1100, 5	2.00	2002	26	8	5.1×10^{-10}	0.099	1.1×10^{-7}
	1100, 10	2.78	2781	15	7	3.2×10^{-10}	0.016	9.4×10^{-8}
	1100, 20	2.14	2140	13	6	5.3×10^{-10}	<0.01	7.9×10^{-8}
	1200, 1	0.487	487	28	11	1.6×10^{-10}	0.76	1.9×10^{-7}

Table S10: (*continued*)

Mixture	T_0 (K), p (bar)	τ (ms)	$\tau/\Delta t$	$\lambda_e \tau$, fit	$\int \lambda_e \, \mathrm{d}t$	Bias for 5%	Delay change for 5% relative error (%)	NO bias for 10%
	1200, 2	0.334	334	33	10	5.1×10^{-10}	0.56	1.7×10^{-7}
	1200, 5	0.305	305	41	9	2.9×10^{-9}	0.26	1.4×10^{-7}
	1200, 10	0.427	427	31	7	5.2×10^{-9}	0.11	1.2×10^{-7}
	1200, 20	0.508	508	18	6	5.0×10^{-9}	0.024	1.1×10^{-7}

Table S11: Hydrogen/air ($\phi = 0.5$): dependence of the ignition tolerances on the step, with the same bias per unit time. Steps: steps of 1 µs in the induction; ratio: bias tolerance per microsecond relative to that at 1 µs.

T_0 (K), p (bar)	Steps	Step (µs)	Delay change for 5% relative error (%)	Bias for 5% per 1 µs	Ratio
1000, 1	201	1	4.22	7.6×10^{-13}	1.00
		0.25	4.39	7.4×10^{-13}	0.97
		0.1	4.42	7.4×10^{-13}	0.97
1100, 5	46	1	3.64	3.0×10^{-11}	1.00
		0.25	4.05	2.8×10^{-11}	0.92
		0.1	4.15	2.7×10^{-11}	0.91
1200, 5	12	1	2.53	6.7×10^{-10}	1.00
		0.25	3.66	4.1×10^{-10}	0.61
		0.1	4.08	3.7×10^{-10}	0.56
1200, 10	11	1	2.52	1.1×10^{-9}	1.00
		0.25	3.40	5.3×10^{-10}	0.48
		0.1	3.95	4.7×10^{-10}	0.43

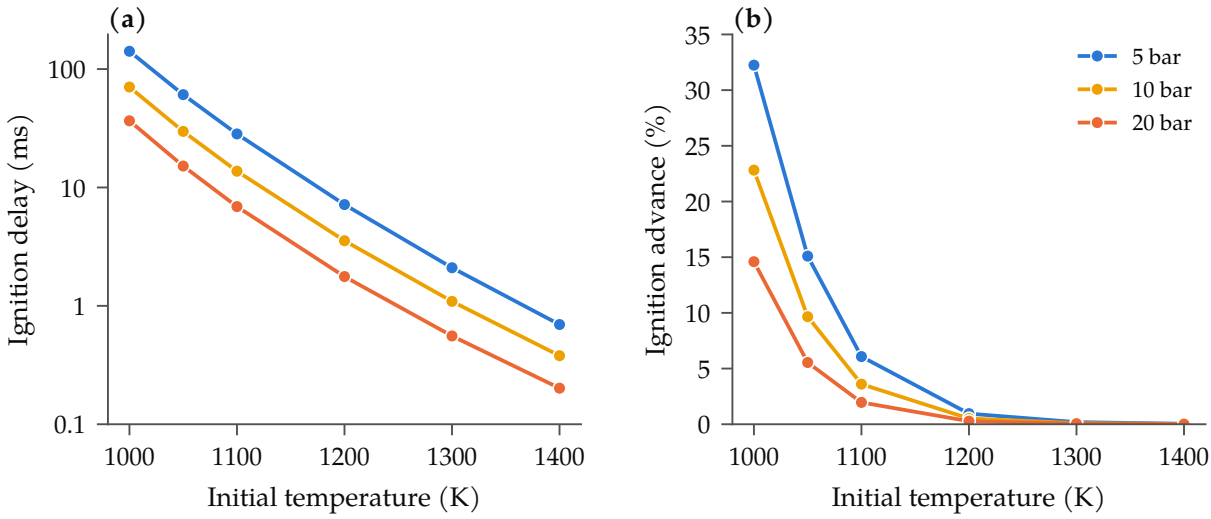


Figure S6: Constant-pressure ignition of methane/air ($\phi = 0.5$, GRI-Mech 3.0) with the exact map at $\Delta t = 1$ µs, without and with a bias of 10^{-11} per step in the mass fractions of H, O, and OH. (**a**) Ignition delay without the bias (time to a temperature rise of 400 K). (**b**) Ignition advance caused by the bias.