Equivalent Frequency Responses for Nonlinear SISO Control-Loop Analysis

TargetDeLab Technical Paper 03

TargetDeLab Group

2026

Abstract

This study examines the use of equivalent linear frequency-domain descriptions for a nonlinear SISO system. The central closed-loop experiment first trains an Attainable Equivalent Frequency Response (AEFR) from the reference-to-output behavior obtained with one specified training waveform. That trained equivalent frequency response is then reused as a linear prediction model for tracking a different validation reference, so that the reader can directly evaluate how well information learned from one closed-loop signal class transfers to another. The spectral description is therefore treated as an experiment-dependent equivalent response rather than a unique property of the nonlinear system. Two distinct descriptions are introduced: (i) the Attainable Equivalent Frequency Response (AEFR), an equivalent closed-loop response conditioned on the controller, training signal, initial state, and constraints; and (ii) the Best Linear Approximation (BLA) of the plant, estimated from repeated random-phase multisine excitations using cross spectra. Both descriptions are used only within their stated validity regions.

The validated methodology is applied to the system $\dot{x} = x^d + u$ with PID gains $K_p = 500$, $K_d = 0.2$, and $K_i = 750$, a $10\,\mu$s time step, and a 1 s time horizon. Different training/tracking signals, identical training and tracking signals, nonlinearity orders d=3,4,5, the constraint $|u| \leq 70$, and variation of the multisine amplitude are examined. The results show that the AEFR can reproduce specific closed-loop conditions very accurately but can be excessively optimistic outside the training class. The BLA of this plant is very close to an integrator over the examined band and shows a small but measurable dependence on amplitude. The study concludes with explicit validity rules for the use of equivalent frequency responses in nonlinear systems.

Keywords: nonlinear systems, frequency response, best linear approximation, BLA, multisine, equivalent linear model, PID, closed loop.

1 Introduction

Transfer functions and frequency responses are fundamental tools for linear time-invariant systems. A general nonlinear system, however, does not possess a unique LTI transfer function that describes all inputs and all trajectories. Nonlinearity generates harmonics, intermodulation, and dependencies on amplitude and operating point; therefore, a simple X/U ratio must be used only as an experimental equivalent under explicitly stated conditions.

The literature distinguishes several families of frequency-domain descriptions for nonlinear systems: describing functions, Volterra/GFRF approaches, nonlinear frequency response functions, and linear approximations in the presence of nonlinear distortion. This study adopts the last direction because it is directly connected to measurements and classical control techniques, while adding explicit validity restrictions.

1.1 Equivalent Linear Models: Purpose and Limitations

Equivalent linearization is useful when the engineering task is prediction or controller analysis over a restricted operating class rather than exact representation of every nonlinear phenomenon. The approximation may be conditioned on an operating point, input amplitude, input spectrum, controller, or probability distribution. The essential scientific requirement is therefore to report the conditioning variables together with the equivalent model. Two experiments performed on the same nonlinear plant can legitimately produce different equivalent linear responses because they solve different approximation problems [1-3,7].

A useful equivalent model must also be validated outside the data used to construct it. Good agreement on the training experiment establishes interpolation capability only for that experimental condition. Validation with a different waveform, amplitude, initial condition, or controller tests whether the approximation transfers to a broader class. This distinction underlies the separation in this chapter between the controller-conditioned AEFR and the open-loop BLA.

1.2 Describing Functions, Volterra Representations, and BLA

Several established nonlinear frequency-domain frameworks address different classes of questions. A describing function retains only the fundamental component of the response to a sinusoidal excitation and is particularly useful for approximate limit-cycle analysis of systems containing static or weakly dynamic nonlinearities. Volterra-series methods generalize convolution by introducing multidimensional kernels; their Fourier transforms are generalized frequency response functions that account for harmonic and intermodulation generation. NARMAX and nonlinear output-frequency-response methods provide alternative data-driven routes to comparable spectral information [3-4,8].

The BLA follows a different philosophy. It does not attempt to reconstruct every nonlinear spectral component; instead, it defines the best linear model for a stated excitation class and treats the remaining nonlinear contribution as distortion. This makes the BLA especially attractive when a linear controller-design workflow is to be retained while explicitly acknowledging nonlinear behavior [1-2]. The AEFR introduced in this chapter is even more conditional because it characterizes a closed-loop behavior that includes the controller and the training trajectory. It must therefore never be interpreted as the open-loop plant transfer function.

1.3 Closed-Loop Identification and the Need for Experimental Conditioning

Closed-loop data create an additional identification issue: the plant input is correlated with disturbances and with the controller action. A simple ratio of Fourier transforms can therefore mix plant dynamics, controller dynamics, reference statistics, and disturbance feedback. Frequency-domain identification methods address this by using appropriate reference or instrumental signals and by stating the feedback configuration explicitly [2,9]. The present simulation environment is deterministic and therefore simpler than a noisy experiment, but the conceptual requirement remains: an identified object must be associated with the experiment that generated it.

For this reason, the chapter uses prediction error on independent scenarios as a primary validity test. An equivalent response is useful when it predicts the closed-loop quantity of interest over the class for which it was constructed. It is not elevated to a universal plant property merely because its Bode-like plot resembles a familiar linear model.

2 Nonlinear System and Control Loop

Topology for equivalent-frequency-response studies. The nonlinear plant is driven by the PID loop for AEFR experiments, with optional amplitude limiting. The AEFR is conditioned on the controller and training reference. The BLA is instead identified from open-loop random-phase multisine excitation and can subsequently replace the plant only as an equivalent linear prediction model.

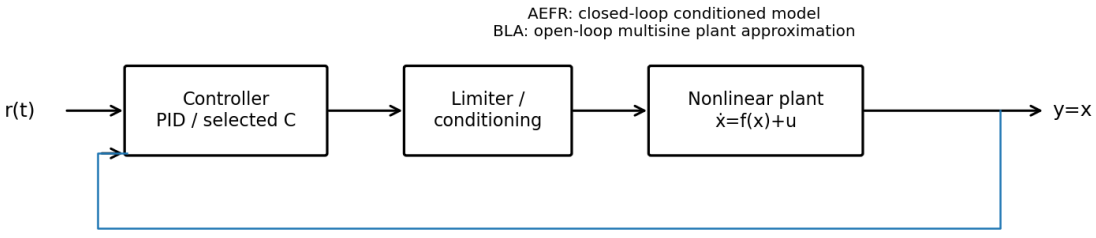


Figure 16. Control topology used in Chapter 3.

$$\dot{x}(t) = f(x(t)) + u(t), \qquad y(t) = x(t).$$

The nonlinear term is represented in the polynomial form

$$f(x) = \sum_{k=0}^{N} a_k x^k.$$

The simulations use the special family $f(x) = x^d$ with output $y = x$. The system is nonlinear for $d > 1$. The initial condition is $x(0) = 0$ unless stated otherwise.

$$e(t) = r(t) - x(t),$$
$$u(t) = K_p e(t) + K_d \dot{e}(t) + K_i \int_0^t e(\tau)\, d\tau.$$

A PID controller with $K_p = 500$, $K_d = 0.2$, and $K_i = 750$ is implemented on the discrete time grid. In saturation scenarios, $\mathrm{sat}_A(u) = \min\{\max(u, -A), A\}$ is used.

2.1 Why X(f)/U(f) Is Not an Exact Nonlinear Transfer Function

In an LTI system, superposition allows G(j ω)=Y(j ω)/U(j ω) to be defined independently of the particular signal, provided the frequency is excited. In a nonlinear system, Y(j ω) may contain frequencies absent from U(j ω), and the response at the same frequency changes with amplitude and operating point. A single spectral ratio therefore does not possess the universality of an LTI transfer function.

2.2 Attainable Equivalent Frequency Response (AEFR)

The Attainable Equivalent Frequency Response (AEFR) is a *closed-loop, experiment-conditioned linear surrogate*. It is not the transfer function of the nonlinear plant. Its purpose is operational: after the nonlinear plant has been operated in closed loop with a specified controller and a specified reference, the observed reference-to-output behavior is summarized as a frequency response over the frequencies that the experiment actually excites.

For a training experiment it is written

$$H_{\mathrm{AE}}(j\omega \mid C, r_{\mathrm{tr}}, x_0, \mathcal{S}) = \frac{S_{xr}(j\omega)}{S_{rr}(j\omega)},$$

where C denotes the controller, r_{tr} the training reference, x_0 the initial state, and $\mathcal{S}$ additional conditions such as saturation. For the deterministic single-record implementation used here,

$$H_{\mathrm{AE}}(j\omega_k) = \frac{X(j\omega_k) R^*(j\omega_k)}{|R(j\omega_k)|^2 + \varepsilon}$$

is evaluated only at sufficiently excited bins. The regularization ε avoids numerical division by zero, but it does not make unexcited bins physically meaningful.

The word *attainable* emphasizes that this surrogate is extracted from behavior that the nonlinear closed loop actually attained under the training conditions. The AEFR may then be replayed as a computationally inexpensive linear predictor for another reference waveform. Such reuse is a validation experiment, not a claim that the nonlinear loop has become LTI.

Training with one reference and tracking another. The identification-and-reuse logic is:

1. close the original nonlinear plant around the specified PID controller and drive it with $r_{\mathrm{tr}}(t)$;
2. record the nonlinear closed-loop output $x_{\mathrm{tr}}(t)$ and estimate H_{AE} from the pair $(r_{\mathrm{tr}}, x_{\mathrm{tr}})$;
3. remove the nonlinear plant from the prediction path and apply a validation reference $r_{\mathrm{val}}(t)$ to the identified AEFR, producing the predicted response $x_{\mathrm{AE}}(t)$;
4. reconstruct the controller action from the same PID law and the predicted error $e_{\mathrm{AE}} = r_{\mathrm{val}} - x_{\mathrm{AE}}$ so that tracking and control-effort metrics can be compared on a common basis; and
5. run the original nonlinear closed loop with the same $r_{\mathrm{val}}(t)$ and compare it with the surrogate prediction.

If the validation waveform has substantially different amplitude, bandwidth, spectral content, or constraint activity, the AEFR can omit nonlinear behavior that was not present in the training experiment. Scenario 1 deliberately tests this situation by training with the accelerated trajectory and validating with the sinusoid. Scenario 2 uses the same sinusoid for training and validation and therefore tests in-class reproduction. The contrast between the two scenarios is central to the interpretation of the chapter.

2.3 Open-Loop Best Linear Approximation (BLA)

For an equivalent linear description of the plant, repeated random-phase multisine excitation is used. The BLA is estimated from the ratio of averaged cross spectra rather than by pointwise division of a single FFT.

$$P_{\mathrm{BLA}}(j\omega_k) = \frac{\sum_m X_m(j\omega_k) U_m^*(j\omega_k)}{\sum_m |U_m(j\omega_k)|^2}.$$

The estimate is valid only at excited lines and depends on the amplitude/phase distribution of the excitation class. Multiple realizations reduce dependence on one particular phase pattern and permit estimation of nonlinear distortion.

2.4 Use of the BLA in an Equivalent Linear Loop

When the controller is linear and saturation is inactive, an equivalent LTI model may be constructed solely for prediction within the BLA validity region:

$L_{\mathrm{BLA}}(\mathrm{j}\,\omega)=\mathrm{C}(\mathrm{j}\,\omega)P_{\mathrm{BLA}}(\mathrm{j}\,\omega)$

$H_{\mathrm{BLA}}(\mathrm{j}\,\omega)=L_{\mathrm{BLA}}(\mathrm{j}\,\omega)/(1+L_{\mathrm{BLA}}(\mathrm{j}\,\omega))$

These relations do not claim that the actual nonlinear system has been converted into an LTI system. They describe the closed loop obtained when the nonlinear plant is replaced by its BLA.

2.5 Pole-Zero Comparisons: Meaningful Scope

A pole-zero comparison is meaningful only after each object has been assigned a legitimate LTI representation. The original nonlinear closed loop does not possess one global set of poles and zeros that is independent of operating point and excitation. Its poles can be discussed only for a stated local linearization or for another explicitly defined linear surrogate. Similarly, the AEFR is measured directly as a closed-loop frequency response; poles and zeros can be extracted only after a finite-order rational model is fitted to the AEFR data. Those fitted poles and zeros then characterize the *trained AEFR surrogate*, not the nonlinear plant. The BLA can be combined with the linear controller and, if represented rationally, yields the poles and zeros of the BLA-based equivalent closed loop.

For the present plant family $\dot{x} = x^d + u$ with $d > 1$, the local linearization at $x_e = 0$ is $P_{\mathrm{lin}}(s) = 1/s$. With

$$C(s) = K_d s + K_p + \frac{K_i}{s}, \qquad (K_p, K_d, K_i) = (500, 0.2, 750),$$

the local closed-loop transfer function is

$$H_{\mathrm{lin}}(s) = \frac{K_d s^2 + K_p s + K_i}{(1 + K_d)s^2 + K_p s + K_i}.$$

The corresponding poles are approximately -415.161 and -1.50544, and the zeros are approximately -2498.50 and -1.50090. These values provide a small-signal reference. A rational fit of the AEFR or the measured BLA can later be compared against this reference, but movement of fitted poles or zeros with the training experiment must be interpreted as variation of the surrogate, not as motion of intrinsic global nonlinear poles.

2.6 Saturation Case

When the amplitude limiter is active, the controller together with the limiter is nonlinear. In that case the algebraic inversion H/(C(1-H)) is not used as a plant estimate. The AEFR is treated directly as a controller-and-saturation-conditioned closed-loop description. This prevents poles and zeros of the equivalent model from being attributed to the actual nonlinear plant.

3 Relation to the Literature

The BLA concept is established in frequency-domain identification of nonlinear systems. Schoukens, Pintelon, Rolain, and Dobrowiecki showed that broadband FRF measurements can be used in the presence of nonlinear distortion, provided that the underlying/best linear approximation is separated from the nonlinear distortion [1]. The book by Pintelon and Schoukens presents BLA, random-phase multisines, closed-loop identification, and model validation in a unified frequency-domain framework [2].

The comparative survey by Rijlaarsdam et al. organizes the main frequency-domain methods for nonlinear systems into Volterra-based methods, nonlinear FRF/Bode approaches, describing functions, and linear approximations in the presence of nonlinearities [3]. The review by Cheng et al. presents Volterra series and generalized frequency response functions as a more rigorous but much more complex alternative for systems satisfying the relevant assumptions [4]. Khalil is used as a standard reference for the general theory of nonlinear systems and local stability [5].

4 Numerical Methodology

The numerical tests use T=1 s, dt=10 microseconds (fs=100 kHz), and x(0)=0. For the BLA, M=100 excited frequencies from 1 to 100 Hz, eight random-phase realizations, and RMS scale 0.10 are used unless stated otherwise. Estimation uses cross-spectral averaging and validity masks. FFT division is never performed at bins with negligible input energy.

For the AEFR, the same basic signals as in the baseline study are used: the accelerated trajectory $r(t) = at^2$ as the training signal in Scenario 1 and the sinusoidal trajectory $r(t) = (a/2)\sin(2\pi\,5t)$ as the tracking signal. Both trajectories are frozen during the second half of the 1 s interval. In Scenario 2, the sinusoid is used for both training and tracking.

4.1 Numerical Consistency Check

The independent numerical implementation reproduced the principal values of the baseline implementation for the nonlinear model. For d=2 without saturation,

IAE=0.05143, RMSE=0.07835, and E_u=1538.56 were obtained, essentially identical to the reference values 0.051427, 0.078343, and 1538.5. For d=5, E_u=5446.36 was obtained, in agreement with the original value 5446.3. This confirms that the independent reproduction preserves the basic plant dynamics or the PID controller.

5 Simulation Scenarios and Results

5.1 Scenario 1 - Different Training and Tracking Signals

Scenario definition. $T = 1$ s, $\Delta t = 10\,\mu$s, $x(0) = 0$, and PID gains $(K_p, K_d, K_i) = (500, 0.2, 750)$ are used. The AEFR is trained with the accelerated trajectory $r(t) = at^2$, frozen during the second half of the record, and is validated on the 5 Hz sinusoidal trajectory $r(t) = (a/2)\sin(2\pi\,5t)$, likewise frozen in the second half. The BLA uses 100 excited frequencies from 1 to 100 Hz, eight random-phase realizations, and RMS scale 0.10.

Table 11. Scenario S1 results.

Scen.	Model	Exit err	IAE	RMSE	E_u / E_u^{rec}	E_{ref}	Limiter %
S1	Nonlinear	0.00182545	0.0514296	0.078345	1538.56	1549.37	0
S1	AEFR	0.00554118	0.00828959	0.0103329	26.9547	1549.37	0
S1	BLA	0.00957771	0.0502681	0.0781431	1529.02	1549.37	0

The AEFR trained with the accelerated trajectory gives very small IAE on this particular sinusoidal test, but the value 26.9547 is not a directly measured actuator energy. It is a reconstructed quantity, denoted here E_u^{rec}, obtained by feeding the AEFR-predicted output into the same discrete PID reconstruction used for comparison. This distinction explains the unusually small value: the AEFR maps reference to predicted output directly and does not enforce the physical state equation $\dot{x} = x^d + u$ during replay. Consequently, the PID action reconstructed from $r - x_{\mathrm{AEFR}}$ is not constrained to equal the physical input that would be required to make the nonlinear plant follow that predicted trajectory. The AEFR is an experiment-conditioned closed-loop equivalent response, not a physical replica of the plant and actuator. When it is driven outside the trajectory class used for identification, the equivalent dynamics can reproduce the output while omitting part of the nonlinear input demand that the true plant requires. The low E_u^{rec} is therefore a surrogate-consistency diagnostic rather than an efficiency gain or a physically measured energy saving. Scenario 3 reinforces the same point: as the nonlinearity order increases, the true plant requires rapidly increasing control action whereas the equivalent linear surrogate remains tied to the excitation regime used for its construction. Thus the energy discrepancy in Scenarios 1 and 3 is rational evidence of loss of validity outside the AEFR identification conditions, not evidence that the physical nonlinear loop can achieve the same trajectory with that small input effort. In contrast, the BLA surrogate remains close to the nonlinear benchmark in IAE, RMSE, and E_u under the tested low-amplitude identification conditions.

The important observation is the mismatch between apparent tracking quality and reconstructed effort. For the nonlinear benchmark, E_u is computed from the actually simulated controller/plant input. For AEFR and BLA replay, the corresponding value is reconstructed from the surrogate output and should be read as E_u^{rec}. Because the AEFR is identified from one closed-loop experiment, it can reproduce selected input-output features while omitting nonlinear mechanisms that determine the actual control action under a different trajectory. The unusually low reconstructed energy is therefore a warning signal rather than a performance advantage. By contrast, the BLA remains close to the nonlinear benchmark in all three reported metrics, indicating better transfer of the low-amplitude linear approximation to the validation experiment.

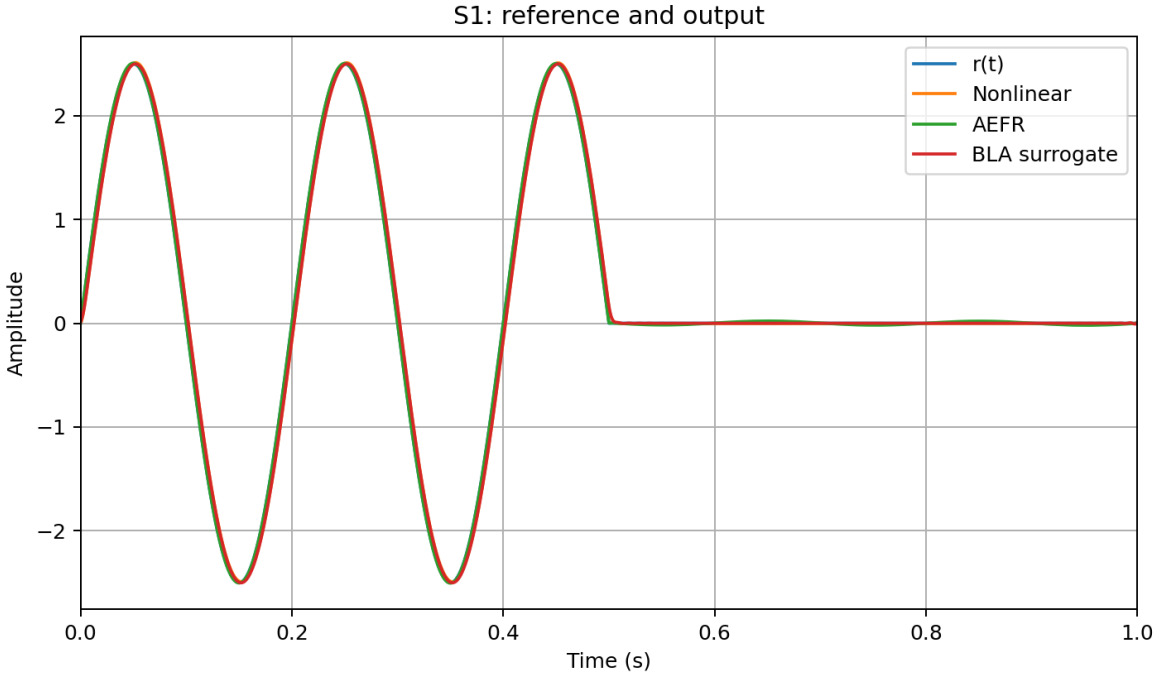


Figure 18. S1 - reference signal and outputs of the nonlinear model, AEFR, and BLA.

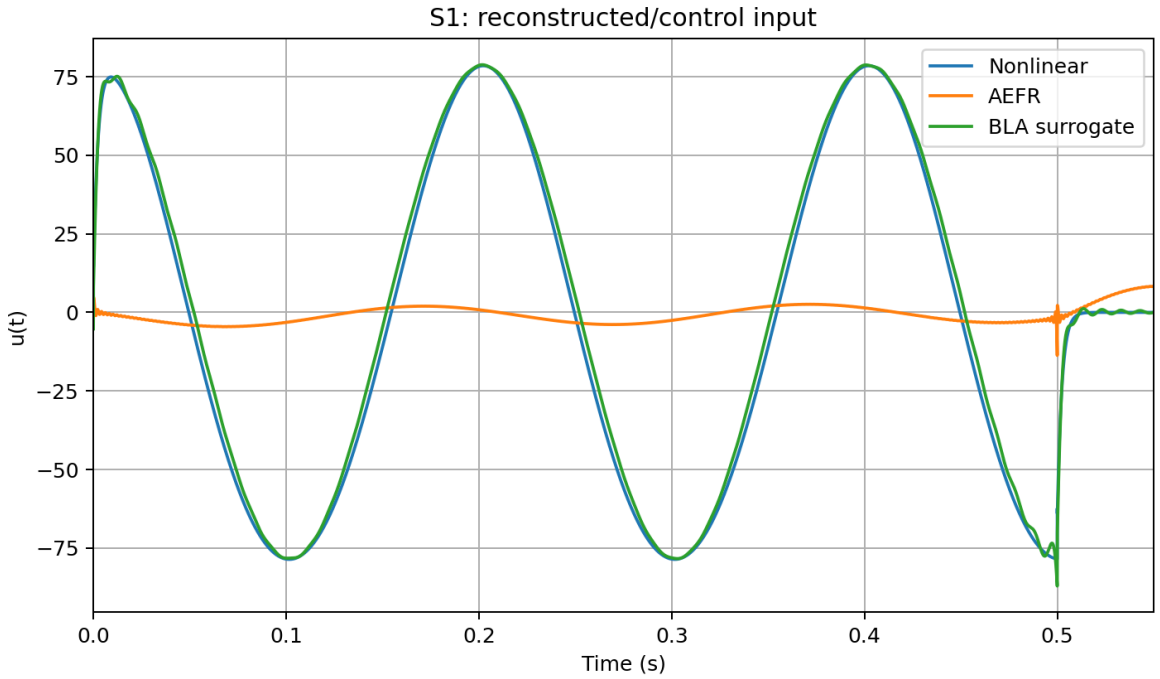


Figure 19. S1 - control/reconstructed-action signals.

5.2 Scenario 2 - Same Training and Tracking Signal

Scenario definition. All numerical and controller settings of Scenario 1 are retained, but the 5 Hz sinusoidal trajectory is used both to train the AEFR and to perform the tracking test. This intentionally tests the in-class fit of the experiment-conditioned equivalent response.

Table 12. Scenario S2 results (verified against the Phase-II MATLAB data).

Scen.	Model	Exit err	IAE	RMSE	E_u	E_{ref}	Limiter %
S2	Nonlinear	0.00182545	0.0514296	0.0783450	1538.56	1549.37	0
S2	AEFR	4.62442×10^{-5}	0.0518721	0.0783043	1535.62	1549.37	0
S2	BLA	0.00957771	0.0502681	0.0781431	1529.02	1549.37	0

When the training and tracking signals coincide, the AEFR reproduces the benchmark very closely: IAE 0.05187 versus 0.05143 and E_u 1535.62 versus 1538.56. This is expected because the AEFR is specifically fitted to the same experiment. The BLA also remains close, with E_u=1529.02.

This near coincidence is an in-sample validation result. It confirms that the AEFR construction can represent the experiment from which it was identified, but it does not establish predictive validity for other references, amplitudes, or nonlinear regimes. The contrast with Scenario 1 is therefore deliberate: a model can fit its training condition extremely well and still become unreliable when the excitation class changes.

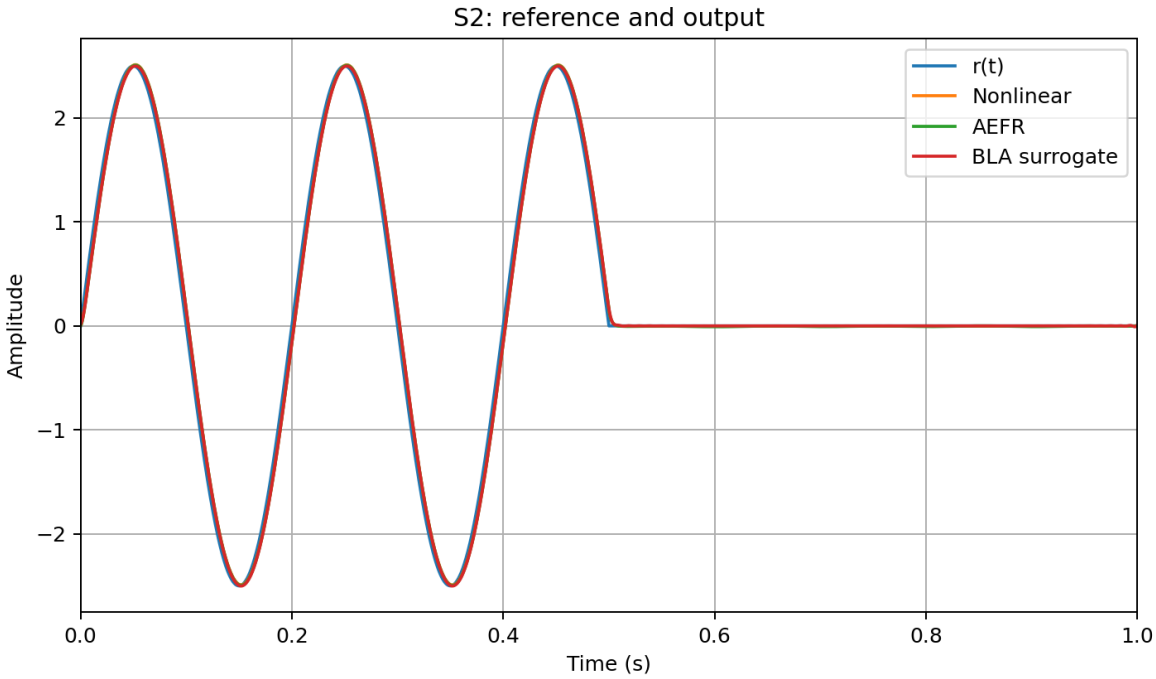


Figure 20. S2 - reference signal and outputs of the nonlinear model, AEFR, and BLA.

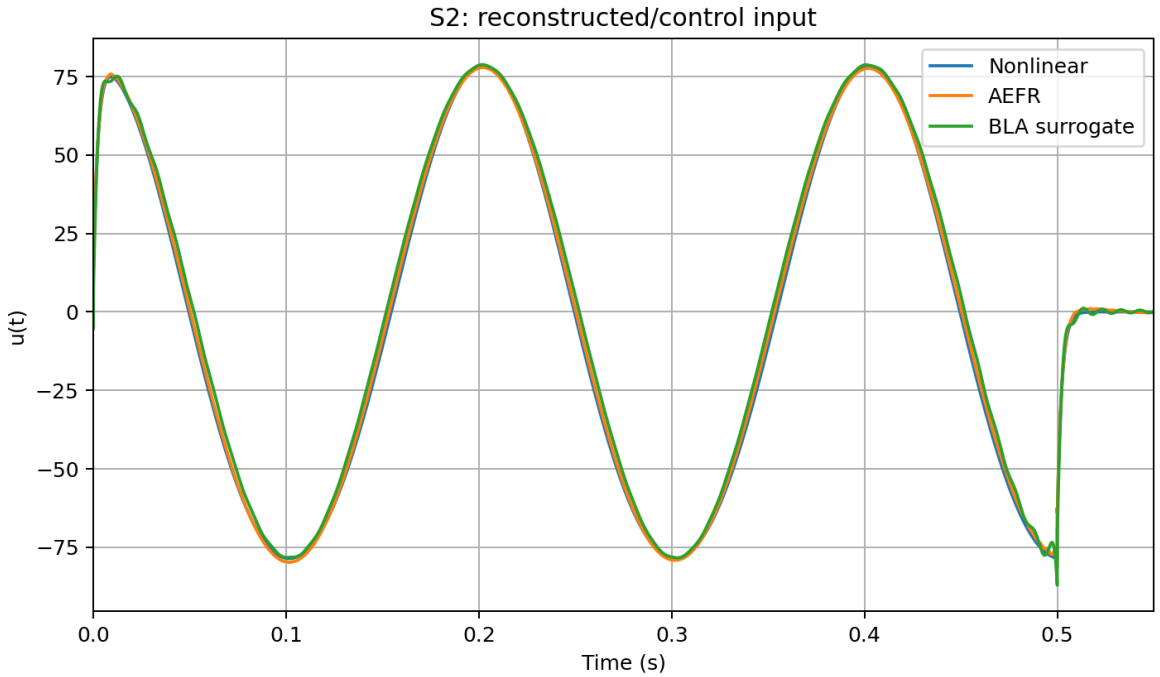


Figure 21. S2 - control/reconstructed-action signals.

5.3 Scenario 3 - Effect of Nonlinearity Order

Scenario definition. The common PID, time grid, and sinusoidal tracking signal are retained while the plant family $f(x) = x^d$ is evaluated for $d = 3, 4, 5$. For each order the

nonlinear benchmark is compared with the AEFR and the low-amplitude BLA under otherwise matched conditions.

Table 13. Results for d=3,4,5.

Scenario	d	Model	IAE	RMSE	$U_E$nergy	E_{ref}
S3-d3	3	Nonlinear	0.0523432	0.0799189	1599.52	1580.2
S3-d3	3	AEFR	0.00877051	0.0108277	29.6104	1580.2
S3-d3	3	BLA	0.0502693	0.0781453	1529.1	1580.2
S3-d4	4	Nonlinear	0.0601831	0.0832591	1774.2	1750.66
S3-d4	4	AEFR	0.00949285	0.0116215	34.1148	1750.66
S3-d4	4	BLA	0.0502693	0.0781454	1529.11	1750.66
S3-d5	5	Nonlinear	0.0944158	0.147291	5446.36	2715.51
S3-d5	5	AEFR	0.010517	0.0127782	41.2379	2715.51
S3-d5	5	BLA	0.0502693	0.0781454	1529.11	2715.51

The actual nonlinear benchmark shows a clear increase in required energy: 1599.52 for d=3, 1774.20 for d=4, and 5446.36 for d=5. The AEFR remains an overly “easy” linear surrogate when trained with the accelerated trajectory, producing extremely small E_u values that must not be interpreted as physically achievable control requirements of the actual plant. This discrepancy is itself a useful diagnostic: the AEFR has left its validity region. Because the specific plant around x approximately 0 has dominant integrator-like behavior, the BLA retains almost the same low-amplitude linear approximation for d=3-5 and likewise does not reproduce the strong high-amplitude nonlinearity of d=5.

As d increases, the growing separation between the nonlinear benchmark and the equivalent linear surrogates quantifies the loss of validity of the linear approximation. The d=5 case is especially informative because the benchmark energy rises sharply while the low-amplitude BLA remains close to its integrator-like form. This is not a contradiction; it reflects the fact that the BLA was identified for a different amplitude regime. The scenario therefore demonstrates why equivalent models must always be accompanied by their identification conditions.

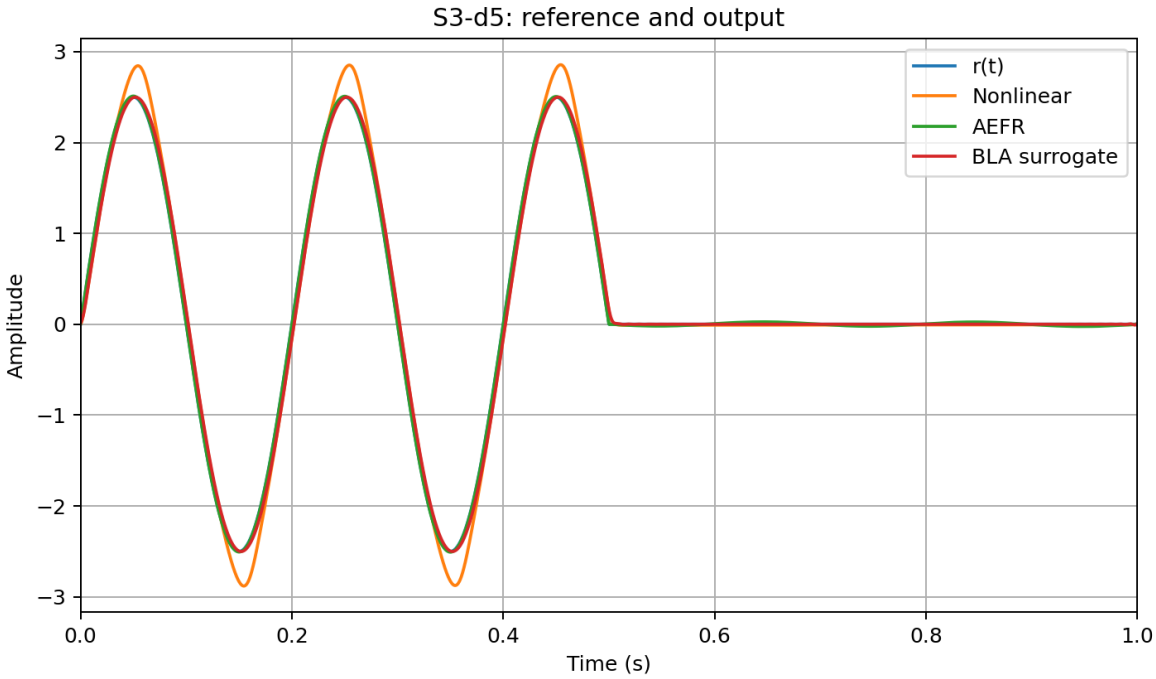


Figure 22. S3-d=5 - departure of the equivalent linear models from the actual nonlinear energy cost.

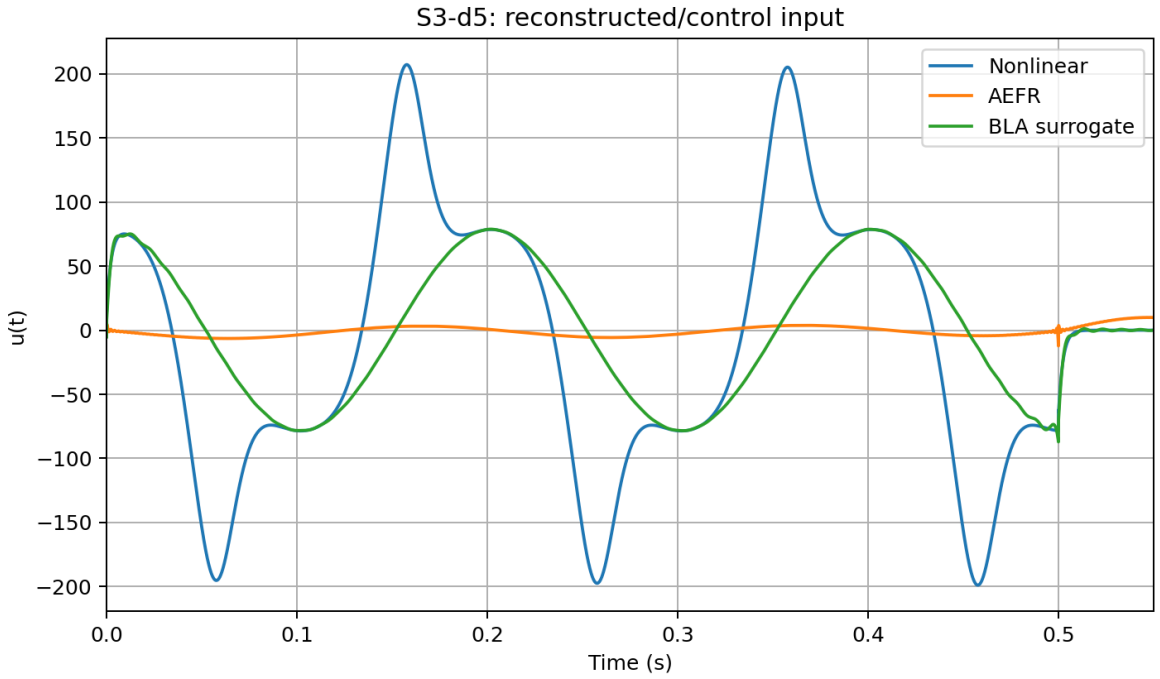


Figure 23. Scenario 3c (d=5) - comparison of control signals.

5.4 Scenario 4 - Amplitude Limiter $|u| \leq 70$

Scenario definition. The d=2 tracking experiment is repeated with the controller output clipped to $|u| \leq 70$. The nonlinear plant, the saturation-conditioned AEFR, and

the BLA-based equivalent prediction are compared. Limiter activation percentage is recorded because saturation makes the controller-plus-limiter itself nonlinear.

Table 14. Results with saturation.

Scen.	Model	Exit err	IAE	RMSE	E_u	E_{ref}	Limiter %
S4-sat	Nonlinear	0.00280304	0.0720308	0.11883	1515.65	1549.37	20.718
S4-sat	AEFR	0.0159735	0.0808071	0.11659	1725.99	1549.37	24.935
S4-sat	BLA	0.00957771	0.0502681	0.0781431	1413.29	1549.37	14.438

The nonlinear benchmark gives E_u=1515.65 with limiter activation for 20.72% of the time, reproducing the baseline study. The AEFR identified under saturation is controller-and-limiter conditioned. Its prediction must not be inverted into a unique plant transfer function. When the reconstructed PID action is clipped at +/-70, the BLA surrogate gives E_u=1413.29 and limiter activation of 14.44%, but this is a prediction of an equivalent model, not proof that the actual nonlinear closed loop will consume this energy.

Once the limiter is active, the closed loop contains an additional nonlinearity that changes both the achieved trajectory and the identification experiment. The different limiter-activation percentages of the nonlinear benchmark and the BLA surrogate show that even a useful equivalent model need not predict constraint engagement exactly. For design purposes, the surrogate can therefore assist comparison, but actuator sizing and saturation analysis must still be verified on the nonlinear closed loop.

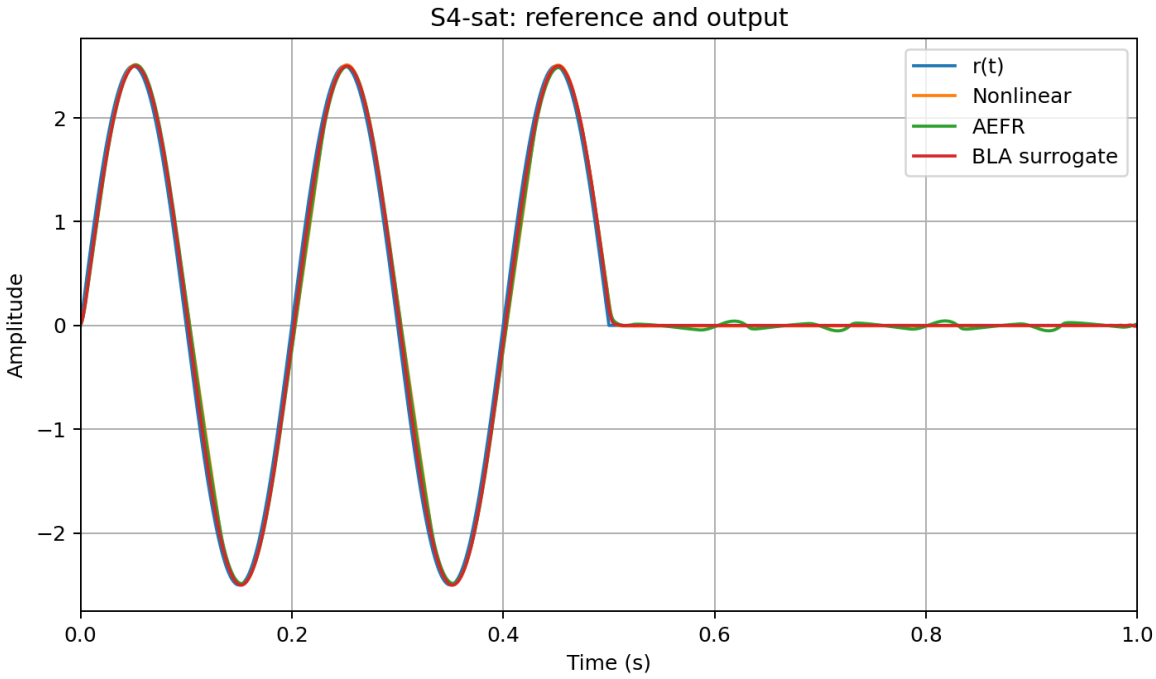


Figure 24. Scenario 4 - tracking with saturation.

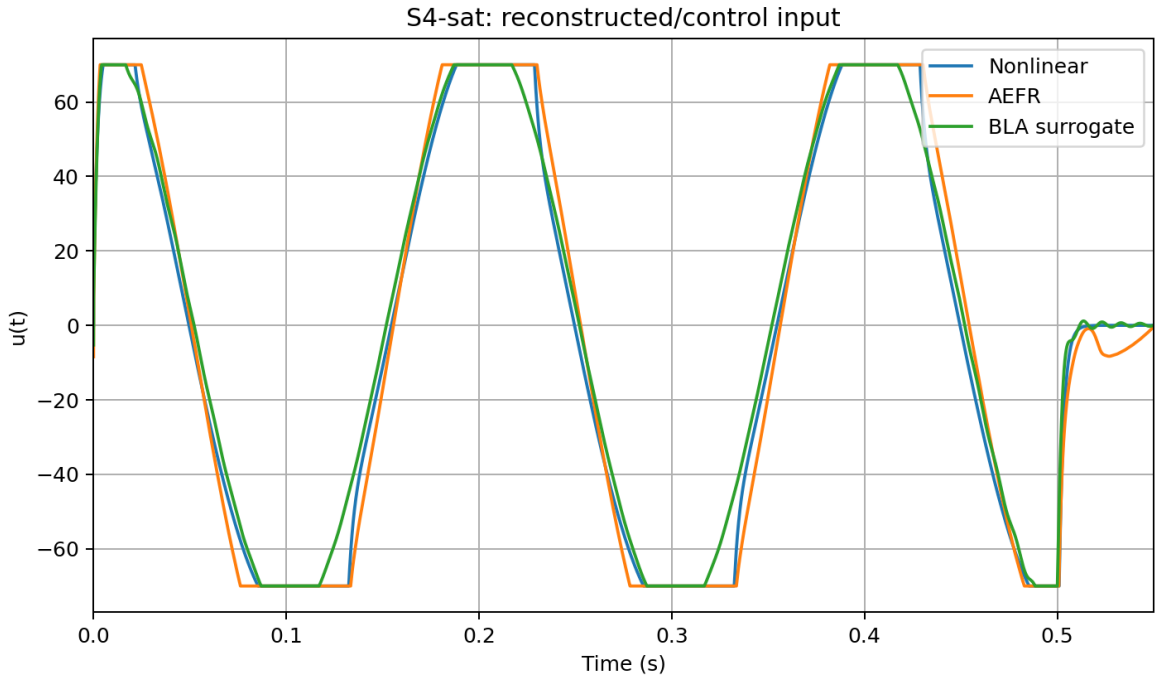


Figure 25. Scenario 4 - effect of saturation on the control signal.

5.5 Scenario 5 - Dependence of the BLA on Amplitude

Scenario definition. Open-loop BLA identification is repeated with the same 1-100 Hz random-phase multisine class and eight realizations, while the RMS scale is changed to 0.05, 0.10, 0.20, and 0.40. The 5 Hz BLA magnitude/phase and nonlinear-output fraction are compared to quantify excitation-level dependence.

Table 15. Variation of the BLA with multisine amplitude.

RMS scale	**$\lvert P_{\mathrm{BLA}}$(5Hz)$\rvert$**	**Phase (deg)**	**NL output fraction**
0.05	0.0318297	-89.9988	1.18252e-09
0.1	0.0318289	-90.0148	2.89481e-09
0.2	0.0318265	-89.9941	1.06605e-08
0.4	0.031847	-90.13	6.10579e-08

At 5 Hz the magnitude remains approximately 0.03183, close to |1/(j2 pi 5)|=0.03183 for an integrator, while the phase is approximately -90 degrees. At the same time, the fraction of nonlinear-distortion energy increases with multisine amplitude. The variation is small in this operating region, but it is methodologically important: the BLA is a property of the system together with the excitation class, not an absolute transfer function.

The small change in the 5 Hz BLA magnitude might suggest amplitude independence if considered alone. The increasing nonlinear-output fraction prevents that interpretation: distortion grows even while the fundamental linear approximation remains nearly unchanged. This is precisely why BLA analysis should report both the fitted linear response and the residual nonlinear contribution. The pair of quantities gives

a more complete indication of whether the surrogate is adequate for the intended excitation level.

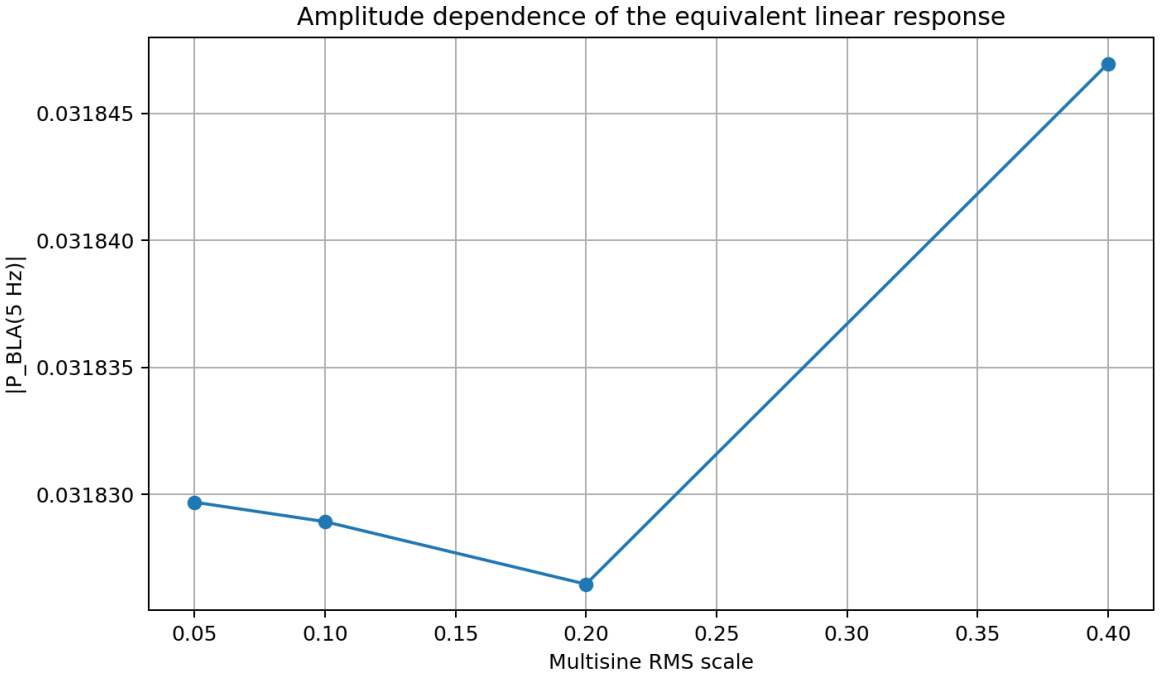


Figure 26. Dependence of $|P_{\mathrm{BLA}}(5\ \mathrm{Hz})|$ on multisine amplitude.

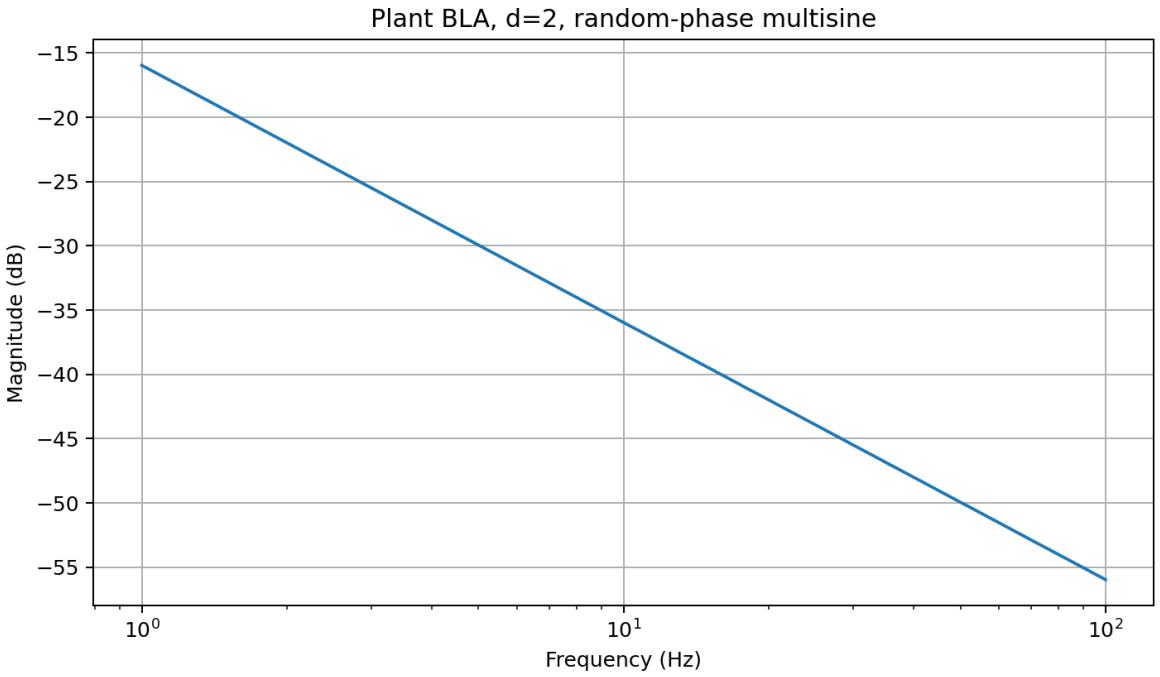


Figure 27. Magnitude of the plant BLA for d=2 and 1-100 Hz excitation.

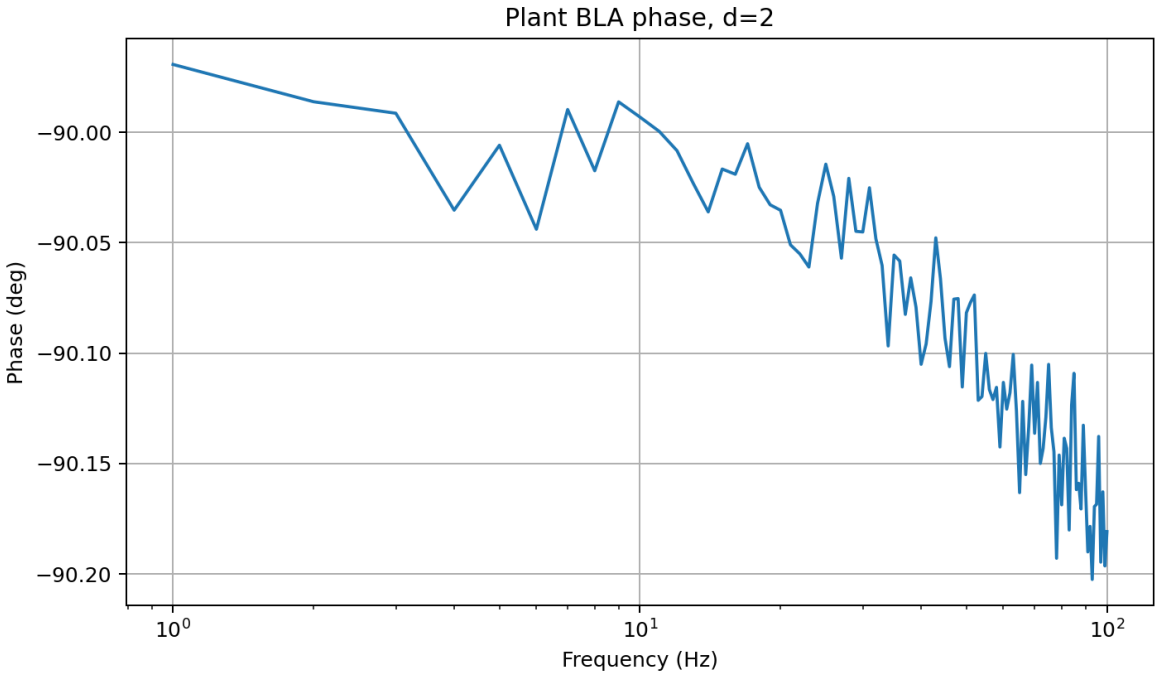


Figure 28. Phase of the plant BLA for d=2.

6 Discussion

6.1 What the AEFR Can and Cannot Claim

The AEFR is particularly useful as a local/data-conditioned projection of the closed loop. It can summarize observed behavior and permit rapid comparison in a frequency-domain form. However, poles, zeros, or stability of the surrogate must not be attributed as intrinsic characteristics of the nonlinear plant. Stability of the actual nonlinear closed loop must be checked by appropriate nonlinear analysis or direct simulation.

6.2 What the BLA Can and Cannot Claim

The BLA is stronger than a simple single-record X/U because it uses ensemble/cross-spectral averaging and explicitly defines the excitation class. It nevertheless remains a linear approximation. For stronger nonlinearities, large amplitudes, or applications in which harmonics/intermodulation are important, Volterra/GFRF, nonlinear state-space, or other nonlinear descriptions are required.

6.3 FFT Validity and Excited Bins

The plant-frequency estimate must use the plant input and output, $P = X/U$, not the reference-to-output ratio. A direct pointwise implementation of $P = X/U$ nev-

ertheless produces artificially large or undefined quantities at bins where the input spectrum U is zero or very small. The implementation uses validity masks and a cross-spectral ratio. Points outside the excited band are not displayed as plant “measurements”.

6.4 Identification and Validation

Using different training and validation signals is essential. Excellent agreement on the training signal itself is not sufficient to prove predictive validity. Scenarios 1 and 3 show why: an AEFR can reproduce one closed-loop experiment yet drastically underestimate energy cost when the trajectory changes or the nonlinearity becomes stronger.

7 Conclusions

The central idea of the baseline study - using frequency-domain descriptions to analyze a nonlinear SISO loop - can be retained provided the interpretation is changed. The ratio X/U from one nonlinear simulation is not an exact or unique transfer function of the plant. It is an experiment-dependent equivalent linear description.

The AEFR is useful when it is stated explicitly that it depends on the controller, training reference, operating region, and nonlinear constraints. The BLA provides a more rigorous open-loop linear surrogate when estimated with random-phase multisines and cross-spectral averaging. In the system examined here, the low-amplitude BLA is very close to an integrator over 1-100 Hz. Increasing nonlinearity and saturation reveal the limits of equivalent LTI models and require validation against the nonlinear state equation.

Across the five scenarios, the main practical lesson is therefore one of validation discipline. Agreement on the training signal demonstrates consistency, whereas agreement on a different signal, amplitude, or constraint condition demonstrates predictive usefulness. The AEFR is most credible inside the experiment class from which it was obtained; the BLA provides a better-defined open-loop approximation but still depends on excitation statistics and level. Neither should replace direct simulation of the nonlinear state equation when strong nonlinearity or saturation becomes important.

With this formulation, the numerical results have a clear scientific interpretation: they do not “prove” that the nonlinear system has one classical transfer function; they quantify how well different equivalent linear frequency-domain descriptions represent specified classes of operation. This is the central contribution of the study.

Relationship to the Complete Book

This paper is a self-contained study extracted and reformatted from the corresponding research chapter of *Nonlinear Control: From Model to Evidence*. The complete book provides the unified progression across all seven studies, cross-study synthesis, extended methodological context, and the MATLAB/Octave and Python software-companion architecture used for replication and further experimentation.

7.1 References

[1] J. Schoukens, R. Pintelon, Y. Rolain, T. Dobrowiecki, “Frequency response function measurements in the presence of nonlinear distortions,” Automatica, vol. 37, no. 6, pp. 939–946, 2001. doi:10.1016/S0005-1098(01)00037-1.

[2] R. Pintelon and J. Schoukens, System Identification: A Frequency Domain Approach, 2nd ed., Wiley-IEEE Press, 2012. doi:10.1002/9781118287422.

[3] D. J. Rijlaarsdam, P. W. J. M. Nuij, J. Schoukens, M. Steinbuch, “A comparative overview of frequency domain methods for nonlinear systems,” Mechatronics, vol. 42, pp. 11–24, 2017. doi:10.1016/j.mechatronics.2016.12.008.

[4] C. M. Cheng, Z. K. Peng, W. M. Zhang, G. Meng, “Volterra-series-based nonlinear system modeling and its engineering applications: A state-of-the-art review,” Mechanical Systems and Signal Processing, vol. 87, pp. 340–364, 2017. doi:10.1016/j.ymssp.2016.10.029.

[5] H. K. Khalil, Nonlinear Systems, 3rd ed., Prentice Hall, 2002.

[6] R. Pintelon, M. Schoukens, J. Lataire, “Best Linear Approximation of Nonlinear Continuous-Time Systems Subject to Process Noise and Operating in Feedback,” arXiv:2004.02579, 2020.

[7] S. A. Billings, Nonlinear System Identification: NARMAX Methods in the Time, Frequency, and Spatio-Temporal Domains, Wiley, 2013. DOI: 10.1002/9781118535561.

[8] A. Gelb and W. E. Vander Velde, Multiple-Input Describing Functions and Nonlinear System Design, McGraw-Hill, 1968.

[9] L. Ljung, System Identification: Theory for the User, 2nd ed., Prentice Hall, 1999.

[10] J. Schoukens, R. Pintelon and Y. Rolain, Mastering System Identification in 100 Exercises, Wiley-IEEE Press, 2012. DOI: 10.1002/9781118218532.