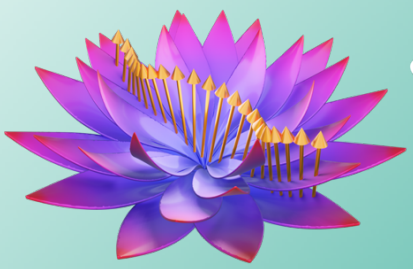

NIMFEIA

Deliverable D6.2

~~Report on demonstration of indoor person recognition with open source data~~

Report on time-series forecasting using magnons

Project Number	101070290
Project name	Nonlinear Magnons for Reservoir Computing in Reciprocal Space
Project acronym	NIMFEIA
Work Package	WP6
Type	Report
Dissemination level	Sensitive
Lead Beneficiary	~~RU~~ HZDR
Due date of delivery	Month 36 – September 2025

Author(s): Zeling Xiong, Christopher Heins, Thibaut Devolder, Fabian Kammerbauer, Mathias Kläui, Helmut Schultheiss, and Katrin Schultheiss

Disclaimer:

The NIMFEIA project has received funding by the European Union's Research and Innovation Programme Horizon Europe under grant agreement No 101070290. However, views and opinions expressed in this document are those of the authors only and do not necessarily reflect those of the European Union. The European Union cannot be held responsible for them.

1. Preamble

Initially, we thought classifying radar data would be an optimal task for our magnetic reservoirs. However, our work on the vortex-based magnon reservoir brought us to the conclusion that working with time-series data instead of radar data is much more suitable. Therefore, we decided to steer our attention to more commonly used tasks for reservoir computing such as time-series forecasting using the Mackey-Glass equation. This time-differential equation is a standard problem to test the performance of physical reservoirs. As a result, the present deliverable reports on how magnons can predict the future of complex time-series instead of reporting on the demonstration of indoor person recognition with open-source data.

2. Introduction

The ability to predict chaotic time series is a longstanding challenge at the crossroads of nonlinear dynamics, machine learning, and physical computing. Chaotic systems, while governed by deterministic rules, exhibit extreme sensitivity to initial conditions, rendering long-term trajectories unpredictable. This combination of determinism and apparent randomness makes them ideal benchmarks for testing new computational architectures aimed at real-time forecasting in complex, high-dimensional systems. Reservoir computing has emerged as a particularly promising approach, offering efficient training and natural suitability for temporal tasks. In this framework, a fixed dynamical system projects inputs into a high-dimensional state space, from which a simple linear readout extracts the relevant information. Crucially, this separation of nonlinear dynamics and trainable output enables hardware realisations, where intrinsic material properties can be harnessed directly. Physical implementations have already been demonstrated in diverse platforms, from optical fibres and memristive circuits to spintronic devices, nanomagnetic arrays, skyrmions, and magnonic systems. Each of these leverages unique mechanisms of nonlinearity and memory.

Here we demonstrate chaotic time-series prediction using a magnon-scattering reservoir (MSR). The platform consists of a $Ni_{81}Fe_{19}$ microdisk in the magnetic vortex state, where magnons provide three key advantages: (i) strong intrinsic nonlinear interactions, (ii) temporal memory through dispersion and damping, and (iii) compatibility with complementary metal-oxide-semiconductor (CMOS) technology, as we have shown in Deliverable 5.1. Harnessing these features, we show that the MSR reliably forecasts the dynamics described by the Mackey-Glass equation, which produces a cyclic yet aperiodic time series widely used to benchmark machine-learning models. Our MSR exceeds previously reported prediction horizons while maintaining high accuracy. This horizon corresponds to approximately six full oscillation cycles of the underlying dynamics. For intuition, these cycles can be viewed analogously to recurrent variations in physiological systems with delayed feedback, such as oscillations in blood cell concentrations governed by the time-delayed response of bone marrow production. Forecasting six cycles ahead in this context would resemble anticipating future phases of over- or underproduction in such a regulatory process. Although purely illustrative, this analogy highlights the temporal depth captured by the magnon reservoir, which reconstructs complex future behavior from the intrinsic memory of its nonlinear dynamics. Our results estab-

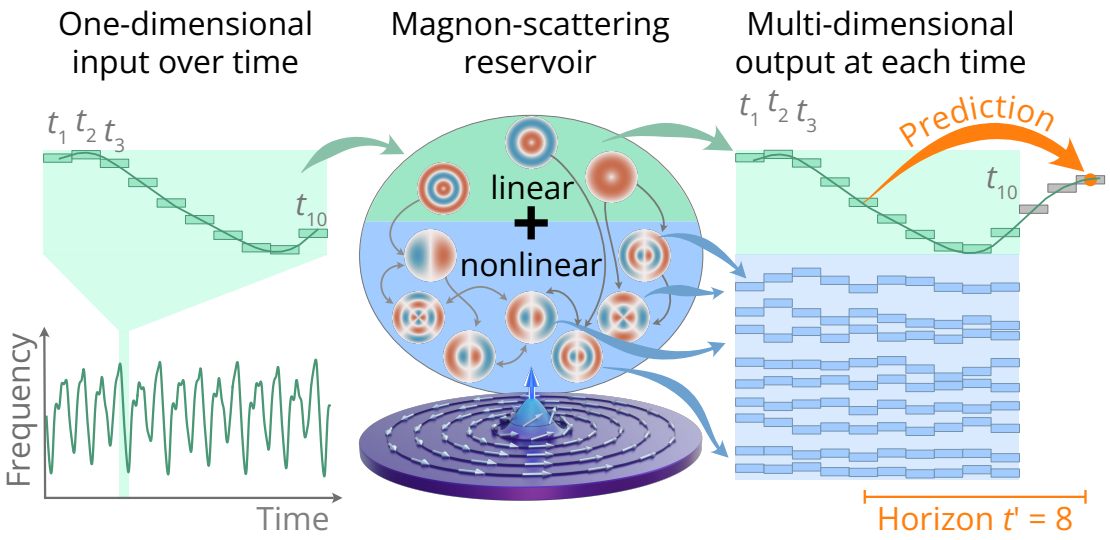

Figure 1. Principle of chaotic time-series prediction based on a magnon-scattering reservoir. A one dimensional temporal input is encoded in a frequency-modulated microwave current and applied to a magnon-scattering reservoir in form of a $Ni_{81}Fe_{19}$ microdisk in the vortex state. Nonlinear interactions between various magnon modes transform the one-dimensional input into a multi-dimensional spectral output trained to forecast the future trajectory of the chaotic signal.

lish magnon-based reservoirs as a scalable physical computing platform for real-time prediction of nonlinear and high-dimensional dynamics.

An overview of the experimental principle is shown in Fig. 1. A time-series input, encoded in a microwave current via frequency modulation, is injected into a micro-antenna for exciting the magnetisation in a $Ni_{81}Fe_{19}$ microdisk. The magnetic ground state is a vortex, where magnetic moments curl in-plane around an out-of-plane vortex core. Nonlinear scattering processes between various magnon modes inside the disk project this one-dimensional, serial input signal into a higher-dimensional spectral output, which is probed via micro-focused Brillouin light scattering (BLS). The resulting data stream is used to train a simple linear readout, enabling accurate prediction of the future trajectory of the input time series.

3. Results

3.1. Mackey-Glass benchmark generation

As a benchmark for our reservoir's prediction capabilities, we use the Mackey-Glass (MG) time series, a prototypical nonlinear model capable of generating chaotic signals, which are widely employed to evaluate reservoir computing systems. The MG sequence is based on a biological process modelling the concentration $P(t)$ of mature red blood cells in living organisms mimicking the time delayed production of new cells in the bone marrow. Essentially, the MG sequence is generated from its delay-differential equation

$$\frac{dP(t)}{dt} = \frac{\beta P(t-\tau)}{1+P(t-\tau)^{\alpha}} - \gamma P(t),$$

where τ is the delay time, β the base level production rate, γ the decay rate, and αthe nonlinearity parameter. Depending on the choice of parameters the equation yields a periodic oscillation of the output value or a chaotic trajectory when the parameters are chosen for strong

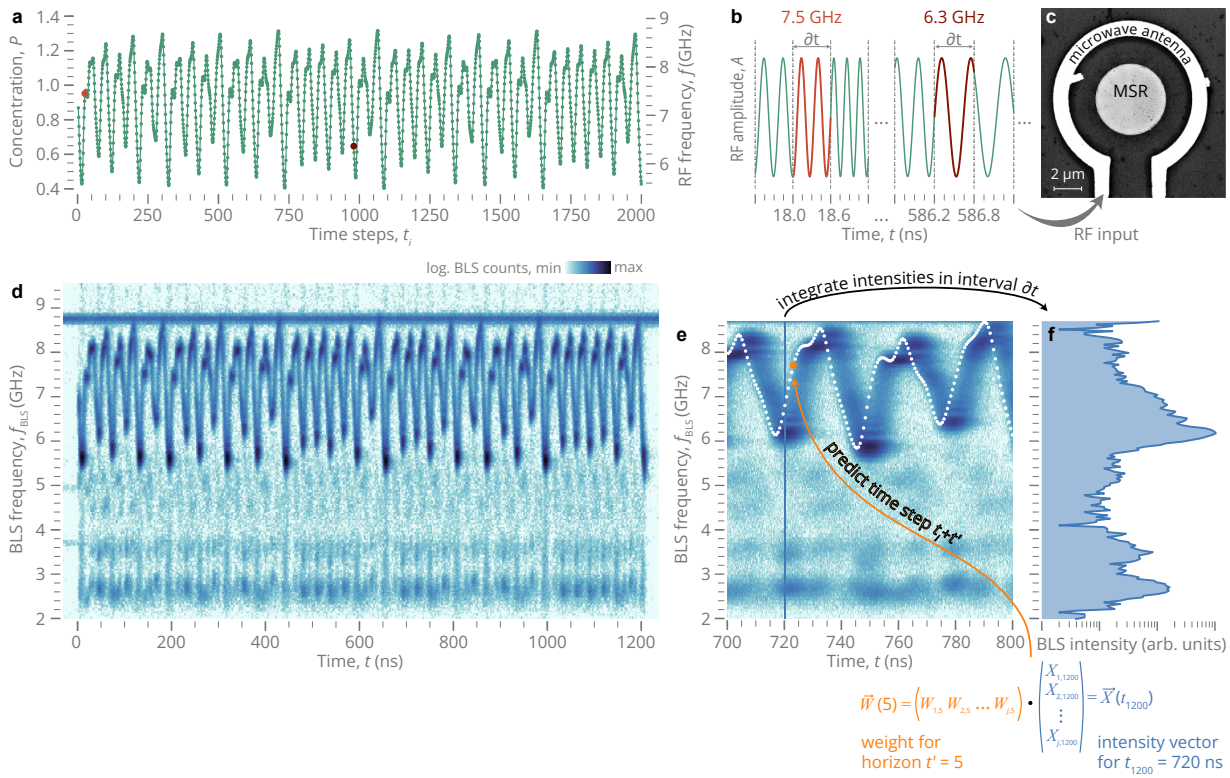

Figure 2. Encoding and reservoir response. (a) The Mackey-Glass time series is mapped onto a frequency-modulated microwave current. (b) Each microwave frequency is held for a certain time window of δt = 0.6 ns and injected into an Ω-shaped antenna surrounding the magnon-scattering reservoir (MSR). (c) Scanning electron microscopy image of the 5 µm-diameter reservoir surrounded by the Ω-shaped antenna for microwave excitation. (d) The time-resolved BLS spectrum reveals both direct responses and nonlinear magnon–magnon scattering, expanding the input into a richer spectral space. (e) Zoomed in section of the data from panel d. The white dotted line indicates the input signals from the arbitrary waveform generator. The measured magnon intensities persist well beyond the direct input window, providing fading memory. (f) The integrated spectral intensities form state vectors that serve as inputs for a linear readout used to generate predictions for a given horizon t'.

nonlinear dynamics. In Fig. 2a, we generated the MG sequence using the parameter set β = 0.2, γ = 0.1, α = 10, and τ = 17, which yields a chaotic trajectory.

3.2. Encoding chaotic input into magnons

To couple this signal into the magnon reservoir, we frequency modulate a microwave current. For each time step t_i the value obtained from the MG equation is linearly mapped $f_{ti} = kP(t_i) + c$ mapped onto a microwave signal. The mapping constants k and c were chosen such that the full dynamic range of the MG sequence matched the frequency interval within the 5.5-8.8 GHz band, as shown by the right y-axis in Fig. 2a. As we know from multiple previous studies, this range of excitation frequencies yields the strongest nonlinear response for a 50 nm-thick and 5 µm-diameter vortex disk. Each discrete frequency value was represented by a sinusoidal burst of duration δt = 0.6 ns (Fig. 2b), synthesised by an arbitrary waveform generator and applied to an Ω-shaped antenna surrounding the MSR (Fig. 2c). To ensure a continuous waveform, the relative phase between consecutive bursts was adjusted during synthesis. A contin-

uous-wave drive at 8.8 GHz, above the nonlinear threshold, was superimposed to increase the number of active magnon modes.

3.3. Time-resolved Brillouin light scattering microscopy

The reservoir's response is probed magneto-optically by time-resolved Brillouin light scattering (BLS). Therefore, A continuous-wave laser (532 nm) was focused onto the sample surface using a 100× objective lens (NA = 0.75), yielding a spatial resolution of approximately 300 nm. The backscattered light was analysed in a Tandem Fabry-Pérot interferometer (TFPI), which resolves the frequency shift due to inelastic photon-magnon scattering with a spectral resolution of 37.5 MHz. Photon counts, interferometer control signals, and a synchronised clock reference from the arbitrary waveform generator were recorded by a time-to-digital converter (Timetagger 20, Swabian Instruments) with 250 ps resolution. By synchronising to the arbitrary waveform generator excitation, the temporal evolution of the magnon spectra was reconstructed. Sample drift was actively compensated using image recognition of in-situ CCD images of the structure, with corrections applied via high-precision piezo stages (Newport XMS series). To access modes with different spatial profiles, µBLS spectra were averaged over 25 positions covering one quarter of the disk (five azimuthal angles at five radial positions). All measurements were performed at room temperature.

Figure 2d shows the spectral intensity during the MG sequence. Within the direct excitation band, the response follows the input frequencies, but crucially, strong signals also appear in the 2-5.5 GHz range due to nonlinear three-magnon splitting. This intrinsic nonlinear mode coupling not only depends on the system's history and provides a unique fingerprint of the dynamic state in which the system is evolving, it also projects the one-dimensional input into a much higher-dimensional spectral output. In addition, the zoomed in spectrum in Fig. 2e highlights that the magnon responses to the microwave excitation signal persist well beyond the 0.6 ns input window, providing fading memory. These two features – nonlinear expansion and memory – are the essential ingredients for reservoir computing. To extract reservoir states, the BLS intensity is integrated within each 0.6 ns time window. Covering the frequency range from 2-8.8 GHz this yields a spectral vector $X(t_i)$ (Fig. 2f) with the number of entries j = 182 determined by the spectral resolution of the BLS measurement. The sequence of vectors forms a state matrix **X** representing the reservoir's transformation of the input. This representation is then used to train a linear readout layer.

3.4. Evaluating reservoir performance

We task the magnon reservoir with predicting the MG signal at $t_i + t'$ time steps into the future. A linear regression model is trained on a subset of 1350 time steps ($51 \leq t_i \leq 1400$, discarding the first 50 steps to allow the system to reach steady operation) to determine the weight vectors $W_{t'}$ for each prediction horizon t':

$$\vec{y}_{\text{training}}(t') = \vec{W}_{t'} \cdot \mathbf{X} + \text{const.}$$

with $\vec{y}_{\text{training}}(t')$ = ($P(51 + t')$,$P(52 + t')$,...,$P(1400 + t')$). The result of this training is plotted in Fig. 3a. The weight vector and constant were found using the *LinearRegression* model from *scikit-*

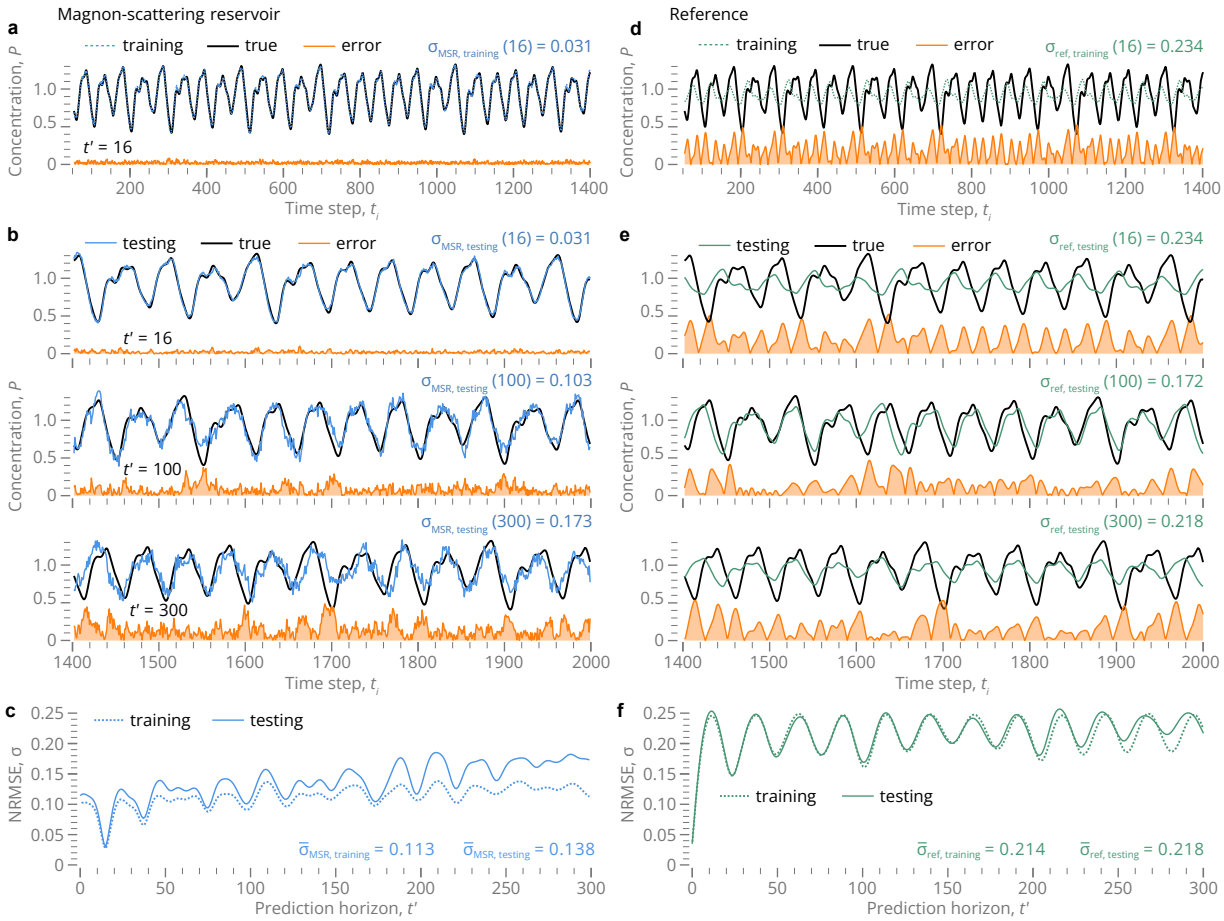

Figure 3. Forecasting chaotic dynamics with and without magnons. (a) Training data for the MSR when predicting t'= 16 steps into the future. (b) Testing results for the MSR for increasing prediction horizon t'= 16, 100, and 300 steps. (c) Normalised root mean squared error (NRMSE) as a function of prediction horizon t', comparing training (dotted) and testing (solid). Viable forecasts are achieved even up to 300 steps, demonstrating competitive long-range prediction performance. (d) Training data of the reference prediction task without MSR when predicting t'= 16 steps into the future. (e) Testing results without the MSR for increasing prediction horizon t'= 16, 100, and 300 steps. (f) NRMSE as a function of prediction horizon t', comparing training (dotted) and testing (solid) for the reference task.

learn, where the model only uses analytical method to find the suitable parameters without any iterative optimisation. This also means that the training time is very short.

Performance is evaluated on the 600 remaining unseen time steps ($1401 \leq t_i \leq 2000$). Figure 3b shows successful forecasts for t' = 16, 100, and 300 steps, with errors increasing gradually with horizon length. Determining an average normalised root mean squared error for each prediction horizon t' (Fig. 3c) shows that the error remains low even predicting 300 steps into the future.

Prediction accuracy was quantified using the root mean squared error (RMSE) between the predicted $\hat{y}_i(t')$ and target $y_i(t')$ values:

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i(t') - \hat{y}_i(t'))^2}\,.$$

To allow comparison across different signal amplitudes, the RMSE was normalised to the dynamic range of the MG sequence, yielding the normalised RMSE (NRMSE):

$$NRMSE = \sigma(t') = \frac{RMSE}{y_{\max} - y_{\min}}.$$

A smaller NRMSE indicates higher predictive accuracy. For each prediction horizon t', the error was computed separately for the training and testing sets to monitor potential overfitting.

To put the MSR's performance into perspective, we follow the same training routine directly on the MG data $P(t_i+t') = w_{t'} P(t_i)$ + const. with weights $w_{t'}$, assuming a scenario without the MSR projecting the time-series input to the higher-dimensional spectral output space. The results are summarised in Fig. 3d-f. Overall, the errors for this reference task remain well above the ones achieved with the MSR across the entire prediction horizon. These results demonstrate that the magnon reservoir achieves competitive forecasting accuracy while directly exploiting nonlinear interactions and fading memory inherent to the magnetic system.

The dimensionality of the reservoir state is determined by the spectral resolution of the BLS output. To evaluate its impact, we applied additional binning to the spectra, effectively reducing the number of features per time step as shown in Fig. 4a. Within each frequency bin, the detected BLS intensity is averaged. Smaller state dimensions simplify the regression task by lowering the number of trainable parameters, which can improve stability during training and mitigate overfitting for limited datasets, as can be seen in Fig. 4b. However, excessive binning discards spectral detail, thereby reducing the richness of the nonlinear mapping and degrading prediction performance. We find that moderate binning yields a favourable compromise between computational efficiency and accuracy, whereas large bin sizes significantly diminish the forecasting horizon. This highlights the importance of tuning the reservoir dimensionality to balance training efficiency and predictive power.

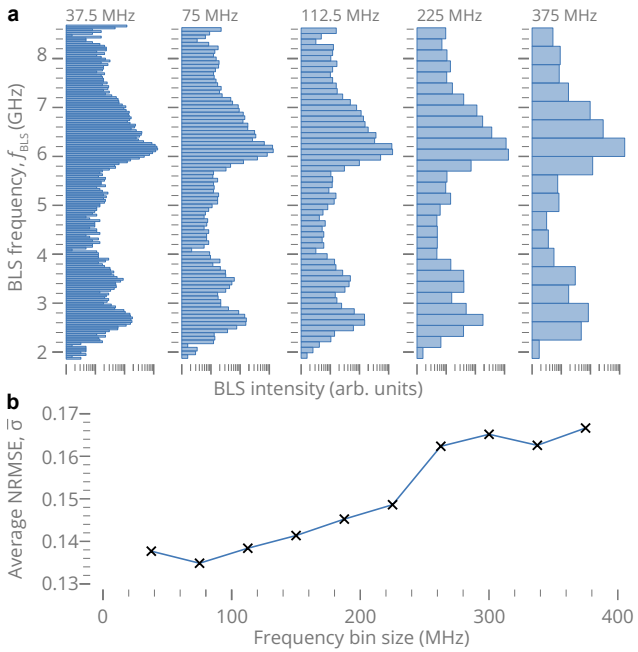

Figure 4. Influence of spectral binning on prediction accuracy. (a) BLS spectra extracted with different frequency bin sizes. (b) Average NRMSE as a function of bin size, showing an optimal performance at twice the natural BLS spectral resolution. Excessive binning reduces spectral richness and degrades forecasting accuracy.

4. Discussion and Outlook

Our results demonstrate that a magnon-scattering reservoir can reliably forecast the Mackey-Glass time series over hundreds of time steps, establishing magnons as a viable substrate for physical reservoir computing. A prediction horizon of t' = 300 corresponds to approximately

six full oscillation cycles of the Mackey-Glass sequence, indicating that the magnon reservoir can reliably capture long-term variations, substantially exceeding previously demonstrated magnetic reservoirs based on artificial spin-ice arrays, which achieved short-term prediction over one to two characteristic cycles with comparable errors. This improvement arises from the continuous, GHz-frequency dynamics of magnons, which provide both stronger nonlinearity and intrinsic fading memory on nanosecond timescales.

Beyond benchmarking on the Mackey-Glass system, the intrinsic properties of magnons suggest several directions for future exploration. First, strong nonlinear interactions and long-lived excitations provide a natural resource for tackling tasks that require both high-dimensional feature spaces and fading memory, such as speech recognition or real-time sensor data processing. Second, the frequency-tunable nature of magnons opens pathways toward multiplexing and parallelism, potentially enabling multi-channel prediction within the same device. Third, compatibility with CMOS processes points toward scalable integration, where arrays of magnetic microdisks could be co-fabricated with conventional electronics.

An additional design principle emerges from our analysis of spectral binning. While reducing dimensionality can simplify training and lower computational cost, excessive compression comes at the expense of predictive accuracy. This trade-off underscores the importance of tailoring the effective state dimension of the reservoir to the specific task at hand. In scalable architectures, such control could be realised by dynamically adjusting spectral resolution or by combining multiple reservoirs with complementary dimensionalities. More broadly, this highlights that not only the physical substrate but also the way its output is encoded plays a crucial role in optimising performance.

From a broader perspective, our work connects nonlinear magnetism with machine learning and neuromorphic computing, demonstrating how fundamental spin-wave physics can be leveraged for tasks traditionally reserved for artificial neural networks. The demonstrated ability to map a simple one-dimensional input into a rich spectral output illustrates how magnonic devices can function as compact, energy-efficient reservoirs. Looking ahead, combining magnon reservoirs with other spintronic or photonic elements may yield hybrid architectures that exploit the strengths of multiple substrates. Such advances could bring physical reservoir computing closer to deployment in edge devices, where real-time prediction and low power consumption are critical.

In summary, the use of magnons for chaotic time-series prediction highlights the potential of magnonics as a platform for unconventional computing. By uniting the fields of nonlinear dynamics, magnetism, and machine learning, this approach opens new avenues for scalable, real-time information processing in both scientific and technological contexts.