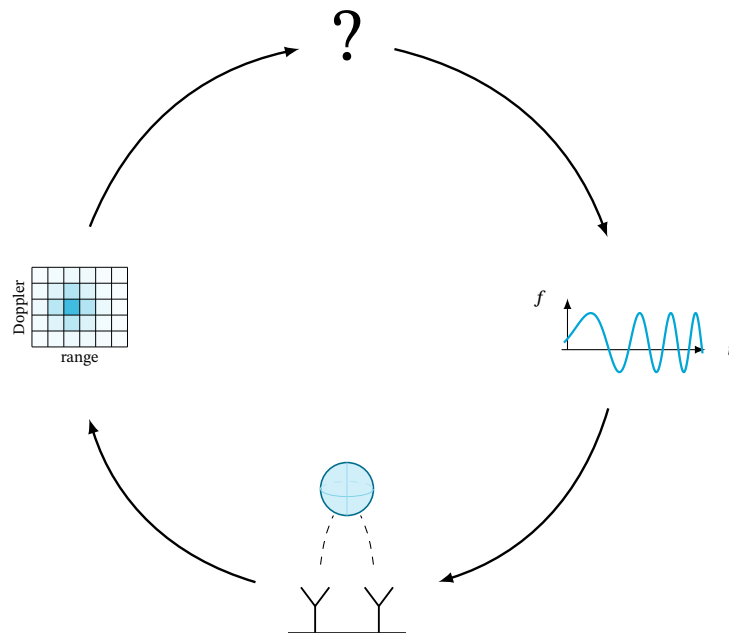


# Research Proposal

Go/No-Go for the PhD thesis:

## Waveform Diversity for Ground Surveillance Radar



by

**Amos HEBB**

Tuesday 16<sup>th</sup> June, 2026

This report has been composed with support from:

Prof. DSc. O. Yaroyvi,  
Dr. F. Fioranelli,  
Dr. J.S. Sandenbergh,  
Prof. Dr. A. De Maio,

Delft University of Technology, promotor  
Delft University of Technology, promotor  
Thales Germany  
University of Naples Federico II

This research was partially funded by Thales Germany through the “Waveform Diversity for Ground Surveillance Radar (WaveRad)” Project.



This research has been carried out at the Delft University of Technology in the Microwave Sensing, Signals, and Systems (MS<sup>3</sup>) group.

Copyright ©2026 by Amos Hebb is licensed under CC-BY 4.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

This report is available at <https://amoshebb.com/doc/hebb2026feasibility.pdf>

Author e-mail: [a.n.hebb@tudelft.nl](mailto:a.n.hebb@tudelft.nl)

# Summary

*Along with designing and selecting waveforms from first principles, this work investigates what can be learned from real hardware returns.*

This PhD asks whether the waveform a ground surveillance radar transmits can be learned from the radar's own returns rather than designed offline. The question matters because the target set has widened to include drones, which are small, slow, and hard to separate from urban clutter, while deployed systems still transmit the Linear Frequency Modulation (LFM) chirps they were designed around. The waveform diversity literature offers many alternatives without converging on any, and evaluates them almost entirely in simulation. To our knowledge, no published ground surveillance radar closes the loop from captured return to next transmitted waveform on real hardware. Closing it would move waveform design from an offline engineering activity to an online, data-driven one.

The instrument is Polarimetric Agile Radar in S-band and X-band (PARSAX), a software-defined dual-polarimetric radar at Delft University of Technology that accepts arbitrary waveforms from a Graphics Processing Unit (GPU)-attached host. On it, 3 experiments form a progression that removes human specification one step at a time. Learned Digital Pre-Distortion (DPD) is handed a target waveform and learns to invert the power amplifier's distortion. Sense and Notch is told what to suppress, a frequency-hopping interferer, and learns from the return alone where to place a spectral notch. Target-Matched Illumination is told only that a drone should be revealed and learns which waveforms reveal it, scored on micro-Doppler signal-to-clutter-plus-noise ratio.

Measurement campaigns in October 2025 (a drone at 500–600 m) and May 2026 (a drone at 9 km) demonstrate dual-polarimetric capture and on-site processing. Programming a campaign's waveforms now takes 1.7 s by script, where the first campaign cost an afternoon at the Arbitrary Waveform Generator (AWG) front panel. In simulation, a single neural network reduces root-mean-square waveform distortion at the amplifier saturation knee by 67–75% for chirps relative to the uncorrected path, and by 51% on a waveform family held out of training, suggesting it learns the amplifier's response rather than a per-waveform correction. European Radar Conference (EuRAD) 2026 has accepted this simulation study for oral presentation. One engineering constraint still blocks rapid iteration. The Field Programmable Gate Array (FPGA) firmware must be rebuilt for each pulse length. Removing that constraint is the first deliverable, and a fixed-pulse-length fallback narrows the waveform search without blocking the core question.

Each experiment ends in a journal submission, with the defense targeted for September 2029. If the loop closes, the question facing this PhD shifts from whether a radar can learn its own waveform to what is worth learning once it does.

# Contents

<b>Summary</b>	<b>3</b>
<b>1 Introduction</b>	<b>10</b>
1.1 Motivation . . . . .	10
1.2 Definitions . . . . .	10
1.2.1 Signal Model . . . . .	11
1.2.2 Pulse Compression . . . . .	14
1.2.3 Ambiguity Shaping . . . . .	14
1.2.4 Windowing . . . . .	15
1.2.5 Doppler Tolerance . . . . .	16
1.2.6 Digital Generation . . . . .	16
1.2.7 Non-Linear Modulation . . . . .	17
<b>2 State of the Art</b>	<b>19</b>
2.1 Waveform Diversity . . . . .	19
2.1.1 Notable Research Groups . . . . .	19
2.1.2 Performance Analysis . . . . .	20
2.1.3 Waveform Design Approaches . . . . .	21
2.2 Ground Surveillance Radar . . . . .	22
2.2.1 Thales Germany . . . . .	22
2.3 Questioning Assumptions . . . . .	23
2.4 Cognitive Radar . . . . .	25
2.5 Digital Pre-Distortion . . . . .	26
2.6 Takeaways . . . . .	26
<b>3 Gaps</b>	<b>28</b>
3.1 Stuck in Simulation . . . . .	28
3.2 Pre-Distortion Mismatch . . . . .	28
3.3 Continual Learning . . . . .	29
3.4 Trustworthy Labels . . . . .	29
3.5 Speculative Directions . . . . .	30
3.5.1 Adapting Only the Transmitter . . . . .	30
3.5.2 Radio-Frequency Tokenization . . . . .	30
3.5.3 Motion-Capture Ground Truth and Multistatic Localization . . . . .	30
3.6 Scope . . . . .	31
<b>4 Feasibility Study</b>	<b>33</b>
4.1 Close the Loop . . . . .	33
4.2 Learned Digital Pre-Distortion . . . . .	35
4.2.1 Problem and Prior Work . . . . .	35
4.2.2 Framework . . . . .	35

---

4.2.3	Simulation . . . . .	36
4.2.4	Hardware Implementation . . . . .	39
4.3	Sense and Notch . . . . .	42
4.3.1	Problem and Prior Work . . . . .	42
4.3.2	Framework . . . . .	42
4.3.3	Simulation . . . . .	43
4.3.4	Hardware Implementation . . . . .	44
4.3.5	Sub-Problems and Extensions . . . . .	44
4.4	Target-Matched Illumination . . . . .	46
4.4.1	Problem and Prior Work . . . . .	46
4.4.2	Framework . . . . .	47
4.4.3	Simulation . . . . .	49
4.4.4	Hardware Implementation . . . . .	49
<b>5</b>	<b>Preliminary PARSAX Experiments</b>	<b>51</b>
5.1	Hardware Chain . . . . .	52
5.2	Preliminary Experiments . . . . .	52
5.3	Second Campaign: Poeldijk, 9 km Baseline (2026) . . . . .	56
5.3.1	Waveform Design . . . . .	56
5.3.2	Tooling . . . . .	58
5.3.3	Results . . . . .	58
5.4	Relation to the Proposed Lines of Work . . . . .	59
<b>6</b>	<b>Administrative Details</b>	<b>61</b>
6.1	Research Plan . . . . .	61
6.2	Doctoral Education Plan . . . . .	62
6.2.1	Transferable (15/15) . . . . .	62
6.2.2	Discipline (0/15) . . . . .	62
6.2.3	Research (>4/15 Completed) . . . . .	63
6.3	Concurrent Activities . . . . .	63
6.4	Data Management Plan . . . . .	63
<b>7</b>	<b>Self-Reflection</b>	<b>64</b>
<b>8</b>	<b>Conclusion</b>	<b>65</b>
	<b>References</b>	<b>66</b>
<b>A</b>	<b>Learned Pre-Distortion of Diverse Radar Waveforms</b>	<b>71</b>

# List of Figures

1.1	2 canonical radar pulse waveforms. Both have the same pulse width $T$ . The chirp achieves finer range resolution by modulating the carrier frequency across bandwidth $B$ rather than narrowing the pulse. . . . .	12
1.2	Spectrograms of the 2 canonical pulse types. Time runs along the horizontal axis, frequency along the vertical. The chirp’s diagonal ridge reflects the linear frequency sweep. The unmodulated pulse’s flat band reflects constant carrier frequency. . . . .	13
1.3	Ambiguity function magnitude (log scale) for the 2 canonical pulse types. Delay $\ell$ runs along the horizontal axis, Doppler shift $f_d$ along the vertical. The chirp sharpens range resolution at the cost of a range–Doppler coupling that must be corrected for moving targets. . . . .	14
1.4	Ambiguity functions of 2 canonical waveforms, both at the same pulse duration $T$ , displayed on the same delay-Doppler axes. Color scale is normalized to the peak in dB, floor clipped at $-30$ dB. . . . .	15
4.1	Current PARSAX hardware interconnect. Dashed callouts (1), (2), and (3) mark the integration gaps. . . . .	34
4.2	Amplitude transfer curves for the hard clipper and the Rapp model at two knee sharpness values, normalized so that $v_{\text{sat}} = 1$ . At $p = 2$ the compression begins well below saturation, which is representative of the PARSAX Gallium nitride (GaN) solid-state amplifier. At $p = 5$ the curve is nearly flat until the knee, then falls sharply, approaching the hard clipper as $p \rightarrow \infty$ . Gradients of the hard clipper are zero at saturation. The Rapp model is smooth everywhere, which is essential for gradient-based training through the simulated amplifier. . . . .	36
4.3	Signal chain for learned pre-distortion. IQWaveNet maps the desired base-band waveform $x_{\text{ref}}$ to a pre-distorted command $x_{\text{cmd}}$ , which the AWG up-samples $8\times$ to a real intermediate-frequency signal at 500 MS/s, amplifies through the Power Amplifier, and the Digital Down-Conversion decimates back to 62.5 MS/s. The dashed bypass shows the uncompensated path to be compared with the proposed approach; $\text{PA}(\cdot)$ denotes the full chain output used throughout this work. The digitized output of this chain is $\text{PA}(x_{\text{ref}})$ , which we use directly as $x_{\text{obs}}$ in the training loop. . . . .	36
4.4	Amplitude and phase error across the pulse duration for one example per waveform family, with and without pre-distortion. The uncorrected baseline (blue) shows a systematic amplitude ramp and phase drift that accumulate over the pulse. The network output (orange) is flat and small throughout. The error pattern is a signature of thermal memory in the amplifier; the network anticipates it rather than reacting to it, requiring the receptive field to span the full waveform duration. . . . .	38

4.5	Hardware-in-the-loop iterative learning control training loop for learned pre-distortion. Solid arrows carry the forward pass; dashed arrows mark stop-gradient paths. The dashed border encloses the physical hardware, which is opaque to the optimizer. . . . .	40
4.6	IQWaveNet task loss under iterative learning control, simulated Rapp amplifier, with no gradient access to the amplifier. Loss falls 87% over 1 kSteps.	41
4.7	World-model training loop for deciding waveforms from raw returns. The encoder converts each return to a latent representation, the world model predicts the next return and reward from that state and a candidate waveform, and the policy trains on rollouts from the model. Solid arrows carry real transmissions, spent only to correct the model, the dashed arrow marks a model rollout. . . . .	43
4.8	Closed loop for learned delay-Doppler shaping. A network reads the range-Doppler map and emits a suppression region $\Omega_{SL}$ marking the cells to clear, the joint random frequency-modulated waveform optimization of Felton et al. [21] renders the waveform set that pushes range sidelobe modulation out of that region, and the radar transmits it. The mask is the agent's action, compact and image-shaped, while producing transmittable constant-modulus waveforms is delegated to the optimizer. . . . .	45
5.1	Signal flow through the PARSAX hardware chain, single-polarization, with the bandwidth-relevant interfaces labeled. A 14-bit integer sample array travels from the host over Ethernet to the AWG, through the power amplifier and a directional coupler that taps a loopback monitor signal, and onward to the transmit antenna. The receive path digitizes at 500 MP/s per polarization channel and returns samples to the host via Peripheral Component Interconnect Express (PCIe), where they are available to the GPU on the same bus. . . . .	53
5.2	Aerial view of the experimental area. PARSAX radar at the origin, chimney target at 1180 m, X-Fields at 600 m. . . . .	54
5.3	Range-Doppler map from the first 2025 PARSAX campaign (50 MHz bandwidth, 0.080 ms sweep time, 25-waveform sweep uploaded manually). The chimney return at 1180 m with near-zero Doppler confirms that the custom waveform was transmitted and returned coherently. This baseline establishes that waveform upload and range-Doppler processing work before any automation is introduced. . . . .	54
5.4	Range-Doppler plot showing a timing anomaly. The ground return and static chimney wander across range bins. The source remains unresolved, candidates include the long-pulse capture settings, the maximum sample rate, the multi-waveform packing scheme, and a newer capture firmware reproduces the same anomaly with the standard 1 ms sweep at a more modest sample rate. . . . .	55
5.5	Spectrogram of Option 1 at 50 MHz bandwidth. Horizontal polarization (H) sweeps 100–150 MHz over 0–1000 $\mu$ s, vertical (V) begins at 1049 $\mu$ s and sweeps the same band to 2049 $\mu$ s. The 2 chirps are non-overlapping by construction, eliminating cross-polarization leakage. The full pulse repetition interval is $2^{20}$ samples (2097 $\mu$ s at 500 MP/s). . . . .	57

- 5.6 Spectrogram of Option 2 at 50 MHz bandwidth, horizontal (top) and vertical (bottom) channels shown separately. Both polarizations sweep 100–150 MHz, the vertical channel starts 80  $\mu\text{s}$  after horizontal, visible as the delayed chirp onset in the lower panel. The stagger places cross-polarization leakage at 12 km, 3 km beyond the 9 km drone hover point. The full pulse repetition interval is  $2^{17}$  samples (262  $\mu\text{s}$  at 500 MP/s), roughly 8 times shorter than Option 1. . . . . 57
- 5.7 HH range-Doppler map of the A13 motorway, 21 May 2026 (50 MHz bandwidth, Option 2, 512-sweep coherent processing interval,  $\approx 134$  ms). *Left*: aerial view aligned to the range axis. *Right*: power relative to the per-bin noise floor; blue lines bound the named scene regions shown at right. Positive velocity is radially away from the radar. Static clutter at 0 m/s is suppressed by slow-time mean subtraction. Some aliasing artifacts have been traced back to now rectified issues with filtering before decimation in receive pipeline. . . . . 59
- 6.1 Planned research timeline (May 2026). DPD runs 8 months overlapping the Setup phase, Sense and Notch and Target-Matched Illumination run 10 months each, thesis 12 months. The gap between J3 and thesis start is intentional: journal revisions and resubmissions land there and fold into the Concurrent row. Blue bars show phases, red diamonds mark key deliverables (Go/No-Go, journal submissions J1–J3, and defense), purple triangles mark conference submission dates: one per experiment row, plus an IEEE Radar Conference submission each fall in the Concurrent row (conferences themselves the following spring). Dashed line is today. . . . . 61

# Acronyms

<b>ADC</b>	Analog-to-Digital Converter
<b>ARL</b>	United States Army Research Laboratory
<b>AWG</b>	Arbitrary Waveform Generator
<b>CFAR</b>	Constant False Alarm Rate
<b>CPI</b>	Coherent Processing Interval
<b>DPD</b>	Digital Pre-Distortion
<b>EuRAD</b>	European Radar Conference
<b>FM</b>	Frequency Modulation
<b>FPGA</b>	Field Programmable Gate Array
<b>GaN</b>	Gallium nitride
<b>GPU</b>	Graphics Processing Unit
<b>IEEE</b>	Institute of Electrical and Electronics Engineers
<b>ISAC</b>	Integrated Sensing and Communication
<b>LFM</b>	Linear Frequency Modulation
<b>LoRa</b>	Long Range, a low-throughput sub-GHz radio protocol
<b>MDP</b>	Markov Decision Process
<b>MIMO</b>	Multiple-Input Multiple-Output
<b>MoD</b>	Netherlands Ministry of Defence
<b>OFDM</b>	Orthogonal Frequency Division Multiplexing
<b>PARSAX</b>	Polarimetric Agile Radar in S-band and X-band
<b>PCIe</b>	Peripheral Component Interconnect Express
<b>PPO</b>	Proximal Policy Optimization
<b>PRF</b>	Pulse Repetition Frequency
<b>RF</b>	Radio Frequency
<b>SAR</b>	Synthetic Aperture Radar
<b>SNR</b>	Signal-to-Noise Ratio
<b>SWaP</b>	Size, Weight, and Power

# 1

## Introduction

### 1.1. Motivation

The dominant paradigm in radar waveform design is physics-first. Given a target model, a noise model, and hardware constraints, classical theory produces a waveform. Decades of waveform diversity research have mapped a far richer design space, polyphase codes, nonlinear frequency modulation, and schemes that adapt to the environment, all analytically well understood but rarely fielded. Deployed systems still mostly transmit pulsed or continuous-wave linear chirps [1]. What is not usually shown in the open literature is a radar that uses its own returns to choose its waveform, searching that design space with learning at a scale closed-form methods are not built for.

The targeted application is ground surveillance radar, where the target set has widened from artillery and vehicles to include personnel and drones, increasingly in cluttered urban settings [2]. Drones are small, slower than conventional aircraft, and generate distinctive micro-Doppler signatures [3], yet detecting them reliably against clutter (returns from ground, buildings, sea surface, or weather rather than the intended target) remains an open problem, especially when they have low or no electromagnetic emissions or act as a swarm.

In order to validate the research ideas developed in this PhD, a radar testbed is needed. In this regard, PARSAX is a software-defined dual-polarimetric radar operated by the MS<sup>3</sup> group at Delft University of Technology whose transmit chain accepts arbitrary waveforms generated by a GPU and captures returns at gigahertz sample rates. The loop between a generated waveform and a captured return is one FPGA firmware modification short of closed: the firmware currently rebuilds per pulse length. This work asks what can be learned with this setup, treating data and compute as abundant.

### 1.2. Definitions

Nearly all papers on waveform diversity modify the emitted signal. Beyond that, no strong consensus emerges about what is and is not considered part of this problem space. In the literature, waveform diversity covers simple LFM parameter changes, pulse coding, frequency agility, polarimetric diversity, transmit beamforming, Multiple-Input Multiple-Output

(MIMO) waveforms, spread-spectrum signaling, Orthogonal Frequency Division Multiplexing (OFDM) radar, mismatched receive filtering, and more. The Institute of Electrical and Electronics Engineers definition does little to narrow things down further.

*Waveform diversity:* Adaptivity of the radar waveform to dynamically optimize the radar performance for the particular scenario and tasks. This may also exploit adaptivity in other domains, including the antenna radiation pattern (both on transmit and receive), time domain, frequency domain, coding domain, and polarization domain.

IEEE Standard for Radar Definitions [4]

The remainder of this chapter establishes the signal and processing models used throughout. The definitions are given for generic waveforms, but with frequent focus on chirp-like frequency modulated continuous waveforms given their prominence in many radars.

### 1.2.1. Signal Model

A radar waveform is a function of time. The physical signal transmitted and received is a real bandpass voltage as for example:

$$x_{\text{rf}}(t) = a(t) \cos(2\pi f_c t + \phi(t)), \quad x_{\text{rf}}(t) \in \mathbb{R}, \quad (1.1)$$

with carrier frequency  $f_c \in \mathbb{R}_{>0}$ , instantaneous amplitude  $a(t) \geq 0$ , and instantaneous phase  $\phi(t) \in \mathbb{R}$ . It is almost always more convenient to strip the carrier and work with the complex baseband representation as:

$$x(t) = a(t) e^{j\phi(t)} = \underbrace{a(t) \cos \phi(t)}_{I(t)} + j \underbrace{a(t) \sin \phi(t)}_{Q(t)} \in \mathbb{C}, \quad (1.2)$$

where  $I(t), Q(t) \in \mathbb{R}$ .  $I(t)$  is the in-phase component, aligned with the carrier.  $Q(t)$  is the quadrature component,  $90^\circ$  shifted from it. The passband signal is recovered as  $x_{\text{rf}}(t) = \text{Re}\{x(t) e^{j2\pi f_c t}\}$ . In hardware, a quadrature mixer extracts  $I$  and  $Q$  by multiplying  $x_{\text{rf}}(t)$  by  $2 \cos(2\pi f_c t)$  and  $-2 \sin(2\pi f_c t)$  respectively and low-pass filtering. Equivalently, multiplying by  $2e^{-j2\pi f_c t}$  and low-pass filtering yields  $x(t) = I(t) + jQ(t)$  directly. Everything from here operates in complex baseband unless real is specified.

In the frequency domain, the Fourier transform pair is:

$$X(f) = \int_{-\infty}^{\infty} x(t) e^{-j2\pi f t} dt, \quad (1.3)$$

$$x(t) = \int_{-\infty}^{\infty} X(f) e^{j2\pi f t} df, \quad (1.4)$$

written  $x(t) \leftrightarrow X(f)$  for brevity. The energy spectral density  $|X(f)|^2$  (J/Hz) describes how signal energy distributes across frequency. By Parseval's theorem,  $\int_{-\infty}^{\infty} |X(f)|^2 df = \int_{-\infty}^{\infty} |x(t)|^2 dt = E$ , the total signal energy.

Unless stated otherwise, all signals are processed as uniformly sampled sequences  $x[n] = x(nT_s)$  at rate  $f_s = 1/T_s$  MP/s (megapoints per second),  $n \in \mathbb{Z}$ . A radar burst, or Coherent Processing Interval (CPI), consists of  $K$  pulses, each of duration  $T$  containing  $N = \lfloor T f_s \rfloor$


fast-time samples. The  $k$ -th transmitted waveform is  $s_k[n]$ , with  $k \in \{0, \dots, K - 1\}$  the slow-time (pulse) index and  $n \in \{0, \dots, N - 1\}$  the fast-time (range) index. When all pulses in a burst are identical we write  $s[n]$ .

Most pertinent to the ground surveillance scenario is *target-matched illumination* where, potentially even pulse-by-pulse, waveforms are crafted to the impulse response of specific targets at specific locations [5, 6].

Waveform parameters include:

- Peak transmit power  $P_t$  (W): energy per pulse is  $E_s = P_t T$ , which sets detection range and SNR.
- Carrier frequency  $f_c$  (equivalently wavelength  $\lambda = c/f_c$ ): governs susceptibility to electronic countermeasures, atmospheric absorption, beam-width, and radial-velocity resolution.
- Pulse repetition interval  $T_r$  (the inverse of the PRF): limits unambiguous range to  $R_u = cT_r/2$  and unambiguous Doppler frequency to  $|f_d| \leq 1/(2T_r)$ .
- Pulse width  $T$ : increasing  $T$  raises  $E_s$  but extends the near-range eclipsed zone during transmission.
- Amplitude envelope  $A[n]$ : constant-envelope preserves power-amplifier efficiency; a shaped envelope can taper spectral sidelobes at the cost of peak power.
- Phase modulation  $\phi[n]$  (equivalently, instantaneous frequency  $f_i[n] = \frac{f_s}{2\pi} \Delta\phi[n]$ , where  $\Delta\phi[n] = \phi[n] - \phi[n - 1]$  is the sample-to-sample phase increment): sets the bandwidth  $B$  and therefore range resolution  $\delta R = c/2B$ , independently of  $T$ .


These parameters interact with each other. A wider pulse increases energy on target but extends the near-range eclipsed zone during transmission, and pulse modulation partially decouples bandwidth from pulse width. Figure 1.1 shows the 2 most basic pulse shapes with key parameters annotated.



**Figure 1.1:** 2 canonical radar pulse waveforms. Both have the same pulse width  $T$ . The chirp achieves finer range resolution by modulating the carrier frequency across bandwidth  $B$  rather than narrowing the pulse.

The energy spectral density  $|X(f)|^2$  makes bandwidth and spectral occupancy directly legible. The Rayleigh criterion tells us 2 targets at the same Doppler are just resolvable in range when separated by  $\delta R = c/2B$ , where  $B$  is the 3 dB bandwidth of  $|X(f)|^2$  and  $c$

is the speed of light. In practice,  $|X(f)|^2$  is estimated by averaging periodograms across overlapping time windows. For waveforms that change instantaneous frequency within the pulse, the spectrogram (a short-time Fourier transform evaluated at successive fast-time steps) shows how frequency evolves across the pulse window, making the linear sweep of a chirp or the hop of an agile waveform (such as change of carrier when using OFDM) directly visible.



(a) Unmodulated pulse: energy concentrated at the carrier frequency throughout the pulse, appearing as a horizontal band.

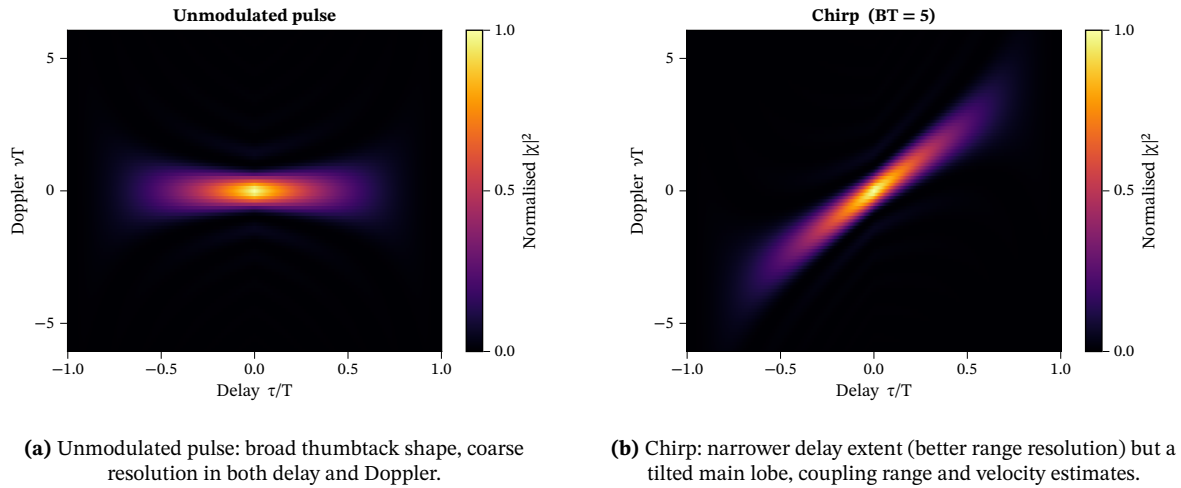
(b) Chirp: instantaneous frequency sweeps linearly from  $f_1$  to  $f_2$ , tracing a diagonal ridge across the time-frequency plane.

**Figure 1.2:** Spectrograms of the 2 canonical pulse types. Time runs along the horizontal axis, frequency along the vertical. The chirp's diagonal ridge reflects the linear frequency sweep. The unmodulated pulse's flat band reflects constant carrier frequency.

The ambiguity function  $A(\ell, f_d)$  is the matched-filter output for a waveform  $x(t)$  as a function of delay (lag)  $\ell$  and Doppler shift  $f_d$  [7, 8],

$$A(\ell, f_d) = \int_{-\infty}^{\infty} x(t) x^*(t - \ell) e^{j2\pi f_d t} dt. \quad (1.5)$$

Its shape makes resolution and sidelobe structure explicit in both dimensions at once. For the unmodulated pulse it has a broad ridge along the delay axis, meaning coarse range resolution. For the chirp the delay narrows (better range resolution) but the main lobe tilts, coupling range and velocity so that a fast target appears displaced in range unless the slope is corrected. Sidelobes outside the main lobe mask nearby targets or generate false alarms. The ideal situation would be the thumbtack ambiguity function, a single impulse at the origin and zero everywhere else. Woodward's volume-invariance theorem [8] explains why no physical waveform achieves this. Formally,  $\iint |A(\ell, f_d)|^2 d\ell df_d = E^2$ , so the integrated volume under the squared ambiguity function equals the square of the signal energy and is invariant to waveform choice. Compressing the main lobe in one dimension redistributes that volume into sidelobes or into the orthogonal dimension. All waveform design is, at some level, a choice of where to push this irreducible volume.



**Figure 1.3:** Ambiguity function magnitude (log scale) for the 2 canonical pulse types. Delay  $\ell$  runs along the horizontal axis, Doppler shift  $f_d$  along the vertical. The chirp sharpens range resolution at the cost of a range–Doppler coupling that must be corrected for moving targets.

### 1.2.2. Pulse Compression

A long pulse focuses more energy on the target, but also means 2 closely spaced targets return overlapping echoes. An unmodulated pulse of width  $T$  resolves 2 targets only if their range separation exceeds  $\delta R = cT/2$ . Frequency modulation breaks this trade-off. Sweeping the carrier across bandwidth  $B$  during the pulse, then correlating the echo against a replica of the transmitted waveform on receive, compresses the received energy into a spike of width  $1/B$  in time, giving  $\delta R = c/2B$  regardless of  $T$ . The time-bandwidth product  $BT$  is the compression ratio: a chirp with  $BT = 100$  provides the range resolution of a pulse 100 times shorter, at the same energy. Pulse width and range resolution are now independent parameters.

The receive filter that maximizes output SNR against additive white Gaussian noise is the matched filter, with impulse response  $h(t) = s^*(-t)$  and frequency response  $H(f) = S^*(f)$  [7]. For a received echo  $r(t) = a s(t - t_0) + w(t)$ , where  $w$  has two-sided power spectral density  $N_0/2$ , the peak output SNR is  $2E_s/N_0$  with  $E_s = \int |s(t)|^2 dt$ . Waveform shape does not appear: only energy sets the SNR floor against white noise. For a single stationary point target, the matched filter output is the autocorrelation of the transmitted waveform. For a linear chirp this is approximately a sinc-shaped pulse, and the first sidelobe sits about  $-13 \text{ dB}^1$  below the peak [7]. Those sidelobes are range sidelobes. For example, a building return 13 dB stronger than a nearby pedestrian will, through its sidelobes, produce a spurious target at the pedestrian’s range that is indistinguishable from the real signal.

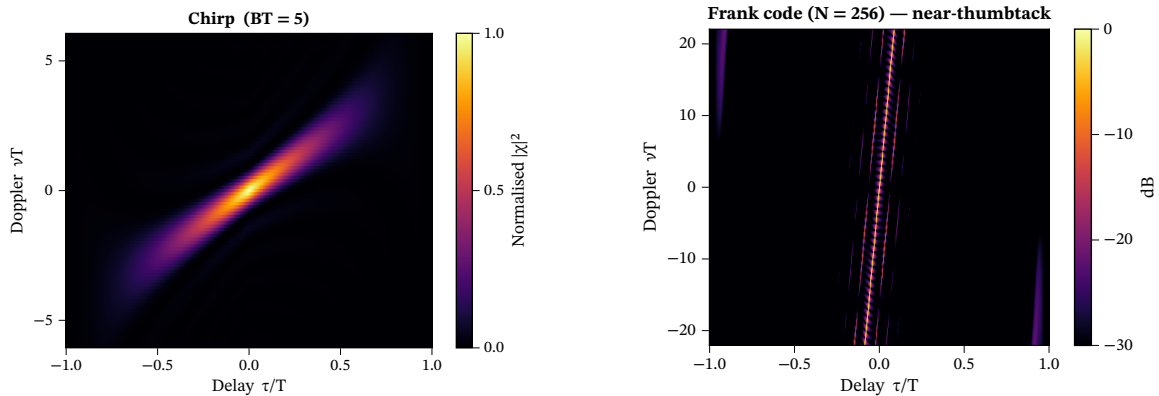
### 1.2.3. Ambiguity Shaping

One of the references for waveform design is *Radar Signals* by Nadav Levanon and Eli Mozeson [7]. Their approach is to treat the ambiguity function not as a side-product of the matched filter but as the primary object to be shaped. Given a desired performance

<sup>1</sup>A decibel is a logarithmic power ratio.  $x \text{ dB}$  corresponds to a factor of  $10^{x/10}$ . So  $-13 \text{ dB}$  is a factor of roughly  $1/20$ ,  $-30 \text{ dB}$  is one-thousandth, and  $+10 \text{ dB}$  is a factor of 10.

profile in delay-Doppler space, the task is to find a waveform whose ambiguity function approximates it as closely as possible, subject to whatever constraints the transmitter imposes. Most waveform design papers in the subsequent literature follow this framing, including most of the entries in table 2.1 later in this report.

Figure 1.4 shows the ambiguity functions of 2 canonical waveforms at the same pulse duration and the same display scale. An unmodulated burst has a wide main lobe in delay: 2 closely-spaced targets cannot be separated in range because the echo of one overlaps the other for nearly the full pulse duration. The linear chirp (figure 1.4a) compresses the delay mainlobe to a fraction of the pulse duration, decoupling range resolution from pulse length, but tilts the main lobe. The tilt means that a slow target at close range and a fast target at greater range can produce nearly identical responses, which complicates tracking in clutter. The Frank polyphase code (figure 1.4b) achieves a near-thumbtack. The main lobe is narrow in both delay and Doppler, and the sidelobe floor sits roughly 24 dB below the peak [7]. Woodward’s volume theorem guarantees the sidelobes cannot be eliminated, only redistributed. The Frank code spreads them as evenly as possible rather than concentrating them along a ridge.



(a) Linear chirp ( $BT = 20$ ). Range resolution improves by a factor of 20, but the tilted ridge couples range and Doppler estimates. (Same chirp ambiguity function as in figure 1.3b, shown here for direct comparison with the Frank code.)

(b) Frank polyphase code ( $N = 256$ ). Near-thumbtack, sidelobes suppressed in both delay and Doppler to roughly  $-24$  dB.

**Figure 1.4:** Ambiguity functions of 2 canonical waveforms, both at the same pulse duration  $T$ , displayed on the same delay-Doppler axes. Color scale is normalized to the peak in dB, floor clipped at  $-30$  dB.

### 1.2.4. Windowing

The  $-13$  dB sidelobe floor of the linear chirp matched filter is often unacceptable for dynamic-range-limited scenarios<sup>2</sup>. Amplitude-weighting the reference waveform before correlating, a technique called tapering or windowing borrowed from spectral analysis, reshapes the sidelobe structure at the cost of a wider main lobe. Concretely, the matched filter reference  $S^*(f)$  is replaced by  $W(f)S^*(f)$ , where  $W(f) \in \mathbb{R}$  is a real, symmetric taper concentrated on the signal band. A smooth  $W(f)$  that rolls off toward the band edges softens the abrupt spectral truncation responsible for the sinc-like sidelobes.

<sup>2</sup>Dynamic range is the ratio of the strongest to weakest signal a system handles simultaneously, without the former corrupting the latter.

Common windows trade mainlobe width against sidelobe level [9]. A Hamming window drops the first sidelobe to  $-43$  dB, widening the main lobe by about 46%. Taylor windows let the designer specify the sidelobe level directly. At  $-35$  dB the main lobe broadens by about 20%. Dolph-Chebyshev windows set every sidelobe to exactly the specified level rather than rolling off.

The tradeoff is not always obvious from the numbers. A 20% resolution loss sounds modest, but in a dense urban environment where 2 targets might be separated by a fraction of a resolution cell, it can mean the difference between resolving them and not. The designer choosing a window is implicitly deciding which matters more, suppressing strong clutter neighbors or preserving fine range discrimination.

### 1.2.5. Doppler Tolerance

A target moving radially at speed  $v$  shifts the frequency of its echo by the Doppler frequency  $f_d = 2v/\lambda = 2vf_c/c$ , where  $\lambda = c/f_c$  is the carrier wavelength and  $c$  is the speed of light. For a linear chirp with slope  $k_r = B/T$ , the Doppler shift  $f_d$  produces a range bias  $cf_d/(2k_r)$  in the compressed output and an amplitude loss  $|\text{sinc}(f_d T)|$  at the filter peak [7]. When  $|f_d T| \ll 1$  the amplitude loss is negligible. The Doppler tolerance of the linear chirp means a waveform designed for stationary targets still performs well against moving ones without modification. This is not a general property of modulated waveforms, it is specific to the linear frequency sweep.

Designing a waveform to have lower sidelobes than the linear chirp typically degrades Doppler tolerance. A nonlinear chirp whose instantaneous frequency profile has been optimized to suppress sidelobes compresses a stationary-target echo cleanly, but a moving-target echo experiences a frequency shift that is no longer compensated uniformly across the pulse. The mismatch between the received chirp trajectory and the designed reference reintroduces sidelobes, sometimes at higher levels than the un-optimized chirp. For a ground surveillance radar whose targets are moving through also moving clutter, this is the decisive practical objection to many theoretically attractive designs.

### 1.2.6. Digital Generation

Chirps can be generated with analog components. Surface-acoustic-wave filters, voltage-controlled oscillators, and dispersive delay lines all produce well-calibrated linear sweeps without digital hardware. These have other advantages, none of which we will explore further. This project only considers digitally generated waveforms from an AWG. A digital generator leaves waveform design as a software parameter, which is flexible enough for exploring anything adaptive.

In a digital system, the matched-filter operation (equation (1.6)) computes pulse compression as:

$$y = \text{DFT}^{-1}\{\text{DFT}\{r\} \cdot \text{DFT}\{s\}^*\}, \quad (1.6)$$

where  $r$  is the digitized received echo,  $s$  is the stored transmitted waveform, and  $*$  denotes complex conjugation. This follows directly from the convolution theorem, where correlation

in time is element-wise multiplication in frequency. The cost drops from  $\mathcal{O}(N^2)$  for a direct time-domain correlation to  $\mathcal{O}(N \log N)$  via the fast Fourier transform.

The frequency-domain form opens the design space in ways that go beyond efficiency. Any modification to the reference spectrum  $\text{DFT}\{s\}^*$  before the inverse transform is equivalent to changing the receive filter, at no additional computational cost. Applying a window function to the reference spectrum corresponds to the time-domain convolution of the reference with the window kernel, but accomplished here by a single element-wise multiply. Replacing the matched reference with one optimized for low sidelobes rather than maximum SNR is equally simple. For pulse-to-pulse waveform diversity, where the reference changes every pulse, this architecture handles it natively. Each pulse uses its transmitted waveform  $s_k[n]$  as the reference, and  $\text{DFT}\{s_k\}$  is computed once per pulse repetition interval  $T_r$ .

Stacking the  $K$  pulse-compressed range profiles and taking a  $K$ -point discrete Fourier transform along the slow-time axis gives the range-Doppler map:

$$Y[n, d] = \sum_{k=0}^{K-1} y_k[n] e^{-j2\pi kd/K}, \quad n \in \{0, \dots, N-1\}, \quad d \in \{0, \dots, K-1\}, \quad (1.7)$$

where  $y_k[n]$  is the pulse-compressed output of the  $k$ -th pulse and  $d$  is the Doppler bin index. A target at radial velocity  $v$  appears at Doppler bin  $d \approx 2Kv f_c T_r / c$ .

### 1.2.7. Non-Linear Modulation

If the matched-filter sidelobe floor is set by the waveform's autocorrelation, the alternative to windowing is redesigning the waveform so its autocorrelation has lower sidelobes in the first place. For  $x(t) = a(t)e^{j\phi(t)}$ , the instantaneous frequency is  $f_i(t) = \frac{1}{2\pi} \frac{d\phi}{dt}$ . A linear chirp sweeps at constant rate  $k_r = B/T$  (the chirp rate), giving  $f_i(t) = f_0 + k_r t$  for  $t \in [0, T)$ . A nonlinear chirp instead chooses  $f_i(t)$  to produce a desired spectral shape. Sweeping faster near the band edges and slower near the center concentrates less energy there and produces a naturally tapered  $|X(f)|^2$ . The autocorrelation sidelobes of the resulting waveform can fall below  $-50$  dB while the waveform remains constant-envelope (instantaneous amplitude  $a(t)$  uniform across the pulse), preserving amplifier efficiency. This is the appeal of the polyphase-coded FM and Pseudo-Random Optimized FM (PRO-FM) families [10, 11].

In practice, the improvement over the linear chirp is harder to realize than the optimized range profiles suggest. The Doppler tolerance argument from the previous section applies with full force. The nonlinear frequency trajectory that produces low sidelobes for a stationary target produces a mismatched response for a moving one, and the suppressed sidelobes re-emerge. A radar whose targets are stationary in range-Doppler space might benefit. One facing drones moving at several meters per second through building clutter will not.

The receive-side complement to waveform optimization is the mismatched filter, which replaces the matched reference with a vector computed to minimize sidelobe energy subject to a mainlobe constraint. For a fixed waveform, the least-squares mismatched filter achieves lower sidelobes than windowing at a smaller signal-to-noise-ratio penalty. For pulse-to-pulse waveform diversity, the regularized minimum mean-square error (RMMSE) approach extends this concept. The filter for the  $k$ -th pulse is  $h_k \propto (C_k + \beta I)^{-1} s_k$ , where  $C_k$  is the estimated interference-plus-noise covariance matrix,  $s_k$  is the  $k$ -th transmitted waveform,

and  $\beta > 0$  is a regularization parameter [12]. The combination of diverse waveforms and adaptive receive processing is currently one of the more promising directions in the literature, though most results remain numerical.

## Organization of this Report

Chapter 2 surveys waveform diversity, ground surveillance radar, cognitive radar, and digital pre-distortion. Chapter 3 identifies the open problems the proposed experiments are designed to address. Chapter 4 proposes 3 experiments, Learned Digital Pre-Distortion, Sense and Notch, and Target-Matched Illumination, built on the shared closed-loop infrastructure they require. Chapter 5 describes the PARSAX hardware chain and 2 preliminary drone measurement campaigns that establish the hardware baseline. Chapter 6 covers the research plan and doctoral education program.

# 2

## State of the Art

This project is concerned with *waveform diversity* and *ground surveillance*. The following describes macro directions of the field of waveform diversity, notable research groups, how the literature judges success, and the design approaches and assumptions that shape it. Ground surveillance constraints are discussed in section 2.2, cognitive radar in section 2.4, and amplifier nonlinearity and digital pre-distortion in section 2.5.

### 2.1. Waveform Diversity

Two broad motivations often appear in the waveform diversity literature: Spectral congestion forcing radars to share bandwidth with other users, and the widening gap between what AWGs can produce and what deployed systems actually transmit [1, 13]. In academia, the main objective function of waveform design is typically lower sidelobes. In industry, the primary consideration is ease of generation, transmission, and reception. Harrington et al. [12] provide a 2024 map of the field covering both waveform design and spectrum sharing. Fulvio Gini surveyed the field in 2021 [14] and produced a list of grand challenges. Table 2.1 at the end of this chapter maps the work discussed here onto the platform and scenario dimensions most relevant to the gaps that follow.

#### 2.1.1. Notable Research Groups

The groups most relevant to this work divide broadly into 4 schools.

The constrained-optimization school treats waveform design as a non-convex program. Given an objective (ambiguity function shape, signal-to-interference-plus-noise ratio, spectral compatibility) and a set of constraints (constant modulus, band limits), find the waveform that satisfies them. Antonio De Maio at the University of Naples Federico II and Nathan Goodman at the University of Oklahoma are the most active contributors [6, 15, 16]. Aubry, Carotenuto, and collaborators at Delft University of Technology produced results directly on PARSAX [17, 18], including a study of amplifier distortion on radar waveforms for spectral coexistence that is direct prior art for Experiment 1 discussed later in this report [19, 20]. The approach produces provably good waveforms when the objective is

correctly specified in advance. This project departs from that formulation by trying to learn the objective from hardware interaction rather than specifying it beforehand.

The random school argues, and demonstrates, that waveform randomness is useful. Shannon Blunt at the University of Kansas Radar Systems Laboratory is characterizing random frequency-modulated waveforms in theory and in hardware [1, 13], Jonathan Owen, Felton, and Ravenscroft have extended this to experimental demonstrations of spectrum sharing and real-time sense-and-notch operation [21–23]. The shared assumption is that a random waveform trades peak sidelobe performance for average sidelobe performance, which is acceptable or beneficial in many operational scenarios.

The machine-learning school typically reaches for a learned component to replace or augment one step in an otherwise standard processing chain. Anthony Martone at United States Army Research Laboratory (ARL) focuses on networked radar, where multiple sensors coordinate and machine learning manages the collective sensing resource [24, 25]. The work is hardware-tested but stops short of closing a fully adaptive waveform loop.

The generative school uses generative models to synthesize waveforms with desired properties rather than optimizing a fixed objective. Visa Koivunen at Aalto University and collaborators have shown that a learned generative model can produce waveforms with controlled ambiguity-function statistics [26]. This is the closest intellectual ancestor of Experiment 1 discussed later in this thesis, though the application there is hardware calibration rather than waveform synthesis.

### 2.1.2. Performance Analysis

Legacy ground surveillance radar was designed for a constrained target set of vehicles and personnel moving at speeds well above the blind-speed threshold, in environments where the radar cross section gap between a vehicle and a bird, for instance, was large enough to filter on. Drones sit in the worst part of that old design space. They are slower than the slow-speed cutoff, smaller than the minimum size threshold, and operating in urban environments surrounded by clutter that can be indistinguishable from the target return.

A glance at drone technology shows them getting quieter. Just a couple examples, the DelFly [27] is a flapping-wing robot, and YCombinator has just funded Ornadyne [28] explicitly to build drones whose flight pattern and radar signature are indistinguishable from those of birds.

Benchmarking against a labeled dataset is not an option. No public large labeled dataset exists for ground surveillance radar, and even if one did, returns are too platform-specific and site-specific. Adding waveform diversity renders the possibility of learning from previous datasets very challenging. The received signal's character changes with the waveform, so a model trained on one waveform family will not generalize to another.

The standard answer is to optimize the transmitted waveform toward a target ambiguity function, shaping the range-Doppler response to push sidelobes away from clutter-heavy regions. This improves detection of weak targets near strong returns, but it requires specifying the right objective before the radar transmits.

### 2.1.3. Waveform Design Approaches

Knowledge-aided approaches use prior information (terrain models, clutter spectra, target impulse responses) to shape the ambiguity function toward desired range-Doppler profiles [15, 16]. There is a large body of work in trying to map the radio-frequency domain to cope with congested environments [17]. Direct constrained optimization treats the problem as a non-convex, often convex-ish, program and produces near-optimal waveforms under the chosen objective, at the cost of specifying that objective in advance [29] while working within a narrow set of waveform designs. Even within a single waveform family and a single objective, the optimization is non-convex and hard. Spanning multiple families, sensors, and competing goals makes it combinatorially worse.

A recent entry into waveform design uses an adversarial framework: a generator produces waveforms that a discriminator network cannot distinguish from the ambient radio-frequency background, while an ambiguity-function-based loss term preserves sensing performance [30]. Ziemann and Metzler report a 90% reduction in radar detectability in simulation under this scheme, a hardware demonstration on PARSAX would make a natural extension to the proposed experiments.

Random frequency-modulation radar waveforms have accumulated experimental support over the last decade. Rather than fix a waveform template, the transmitter generates pulses whose instantaneous frequency varies pseudo-randomly pulse to pulse within a spectral envelope. Constant envelope is preserved throughout, making these waveforms compatible with saturated power amplifiers. The pulse-to-pulse diversity prevents range sidelobes from accumulating coherently across the Doppler processing interval, giving this family a structural advantage over fixed waveforms in clutter-heavy scenarios. Blunt et al. [13] provide the principles treatment, Owen et al. [31] give the analytical framework for spectral efficiency.

Linear chirps are recognizable. A receiver tuned to the right band can detect the radar, match the waveform template, and exploit the information. A noise-like waveform, whose autocorrelation mimics white noise within the pulse duration, gives nothing away. A linear chirp produces a diagonal ridge in the delay-Doppler plane that confounds slow targets at close range with fast targets at greater range. A noise-like waveform approximates a thumbtack, with low sidelobes in both dimensions. For cluttered ground surveillance, both properties matter.

The case for noise or communication like waveforms is not theoretical only. Ravenscroft et al. characterized the trade-space experimentally for waveforms that combine spectral notching with noise-like sidelobe structure [23]. Owen et al. demonstrated a real-time open-air sense-and-notch radar running random FM waveforms [22]. More recently, their group has shown that random FM waveforms can be spectrally shaped to approach optimum range sidelobe levels and that the delay-Doppler ambiguity can be jointly optimized across a pulse set [21, 31]. The design toolbox is mature and the real-time implementation has been shown to work.

Generating a spectrally shaped random FM waveform requires solving an optimization for each new pulse. Constant Envelope OFDM sidesteps this. Standard OFDM modulates independent subcarriers and has a high peak-to-average power ratio, making it incompatible

with saturated amplifiers. Applying phase modulation to the OFDM baseband signal maps it to constant envelope. Biehl et al. [32] show that the result is a member of the random FM family: spectral shaping comes for free from the OFDM subcarrier structure, without per-pulse optimization, at computational cost low enough for real-time hardware. Kramer et al. [33] extend this to a non-repeating continuous-wave waveform with a compact parameterization, with open-air measurements.

Beyond mitigating sidelobes, the ambiguity function's two-dimensional structure offers a richer objective. Ground clutter concentrates in specific delay-Doppler cells, with stationary buildings at close range and slow-moving foliage in a band near zero Doppler. A waveform set optimized to suppress sensitivity in those cells while preserving it elsewhere is strictly more useful than one that minimizes sidelobes uniformly. Felton, Owen, and Blunt [21] demonstrate this directly, jointly optimizing sets of random Frequency Modulation (FM) waveforms to suppress range sidelobe modulation in a specified region of the delay-Doppler response, the Target-Matched Illumination experiment builds on this framework. The over-the-air results at 3350 MHz show well-shaped point-spread functions. The same framework admits a cognitive capability. The system senses the delay-Doppler distribution of clutter and interference, then adapts the waveform set to avoid it. These are among the jumping off points for *sense-then-notch* ideas described in section 4.3.

## 2.2. Ground Surveillance Radar

The open literature on ground surveillance radar lags operational capabilities. Publicly available information on deployed systems is limited to marketing materials that do not document known shortcomings. The project sponsor's primary concern is *dynamic range* and, later in our conversations, *low probability of intercept*. They also mention mostly using *pulse frequency modulated* waveforms, and having little interest in OFDM or random waveforms, despite this being somewhat at odds with the latter goal.

### 2.2.1. Thales Germany

What is already in the field sets a floor for novelty. The products below are operational Thales ground surveillance systems. Anything genuinely new either has to exceed them or reach comparable performance.

The following radar models were listed on the Thales Germany radar systems page.<sup>1</sup>

- **Ground Observer 12 (GO12)**: Medium-range tactical radar.  $360^\circ$ ,  $r \leq 27$  km,  $v \leq 125$  ms<sup>-1</sup>. Agile frequency mode. 12–18 GHz. Track-while-scan up to 40 targets.
- **Ground Observer 20 MM (GO20-MM)**: 3D,  $360^\circ$ , 0.01 m<sup>2</sup> drone at 4 km, 0.1 m<sup>2</sup> drone at 6 km, 0.5 m<sup>2</sup> person at 9 km, 2 m<sup>2</sup> car at 13 km, 5 m<sup>2</sup> Helicopter or 10 m<sup>2</sup> trucks limited by 15 km instrumented range.
- **Ground Observer 80 (GO80)**: Successor to “BOR-A 550/560”,  $360^\circ$ , 0.5 m<sup>2</sup> person at 24 km, 20 m<sup>2</sup> vehicle at 60 km, up to 80 km on “larger targets”. 8–12 GHz. Shown with a single exposed parabolic antenna (two different elliptic geometries shown

---

<sup>1</sup><https://www.thalesgroup.com/en/europe/germany/radar-systems>, accessed 2026-02.

online) that do not appear to rotate fast, 1.4° beamwidth. Internal *track-while-scan* up to 50 targets with *Automatic Target Classification* function.

- **Ground Alerter 10 (GA10):** 360°, up to 10 km range, “Counter Rocket, Artillery, Mortar”

Direct quotation withheld at sponsor request, the following themes recurred in multiple conversations:

- **Dynamic range** is the primary operational concern. The example given most often: detecting a pedestrian in the range-Doppler bin of a large building, or a drone among automotive traffic.
- **Urban ISR** as a growth application. “High dynamic range” scenarios in urban environments were described repeatedly.
- **Pulsed Doppler** preferred over continuous wave. This was stated as a preference, not a requirement.
- **Amplifier uniqueness.** Each high-power amplifier is different, calibration is time-consuming and expensive. This is one reason DPD techniques are interesting to them.
- **Advanced architectures.** Thales expressed interest in more software-defined platforms, not just incremental improvements to existing hardware.
- Thales mentioned rain as a nuisance clutter source.
- They raised low-probability-of-intercept in later conversations, after initially dismissing noise radar.

Open-source reporting on recent operational deployments is sparse. The Brave1 marketplace<sup>2</sup>, a Ukrainian defense innovation platform, hosts a dataroom of battlefield sensor data including radar signatures and drone flight patterns, but access requires Ukrainian Ministry of Defence authentication and the data is not publicly available. Requests for access as a PhD candidate have gone, appropriately so, un-answered as of Spring 2026. Press coverage confirms the dataroom exists and contains radar-relevant data. Informal discussion with Canadian artillery officers stationed in Lithuania mention integration, in particular with acoustic systems, and general reliability as issues. Neither are citable sources for this work.

## 2.3. Questioning Assumptions

Several design assumptions in the radar literature are outdated or imposed by constraints that have quietly stopped binding. The following are the ones most relevant to this project.

**Compute and data are not scarce.** The radar literature declares techniques intractable without asymptotic “big  $\mathcal{O}$ ” analysis. This makes distinguishing between genuinely hard problems and ones that are merely “very slow as a naive MATLAB script” difficult. For example, brute-forcing a search over poly-phase codes may be hand-waved away as intractable, despite only having millions or billions of candidates. At

---

<sup>2</sup><https://brave1.gov.ua>

$\mathcal{O}(N \log N)$  per evaluation, a modern GPU clears this in milliseconds, fast enough for real-time waveform selection per pulse. Size, Weight, and Power (SWaP) arguments and appeals to dataset scarcity then justify any reduction of computation or sample efficiency without asking which category applies. This work treats both compute and data as abundant. Consumer grade GPUs are crossing terabytes per second of FFTs, hobbyist training datasets are reaching petabyte scale on home PCs, and Analog-to-Digital Converters (ADCs) are creeping toward direct Radio Frequency (RF) sampling speeds.

**The processing pipeline may change.** The field typically treats the standard chain between the ADC and any downstream processing. Hand-crafted feature pipelines once compressed pixels into something a classifier could handle. Convolutional networks then learned features end-to-end from raw pixels [34], and the hand-crafted pipelines lost. The underlying physics feels too useful to ignore entirely, but the range-Doppler map discards phase and temporal structure that does not fit a coherent processing interval, and its assumptions (linearity, stationarity, a known noise floor) are known simplifications.

**Ground clutter is highly structured.** The dominant returns in a ground surveillance scenario are not random noise. Buildings, roads, and terrain produce strong returns at predictable locations in the range-Doppler map. The matched filter, built on a Gaussian noise model, treats all of this as undifferentiated interference and discards the prior knowledge the structure provides. The structure is also an opportunity. Terrain databases and repeated observations carry prior knowledge about where strong returns will appear. Knowledge-aided waveform design exploits this directly, shaping the ambiguity function to suppress sensitivity in clutter-heavy delay-Doppler cells [16, 17].

**Polarization carries information.** The radar cross section of a target depends on the polarization of the illuminating signal, not just the target's size and shape, and scattering can rotate the polarization of the return. A polarimetric radar, which transmits and receives on two orthogonal polarizations simultaneously, uses this to distinguish targets that look similar in range and Doppler.

**Spectrum sharing and intercept probability are linked.** Ground surveillance radars increasingly share spectrum with commercial users. Waveforms that can place notches around occupied bands without sacrificing range resolution may become necessary. Adversarial environments add a second constraint. A waveform an adversary can intercept and use to predict the radar's next transmission is a liability. Random FM and noise-like waveforms have low probability of intercept as there is no template for an adversary to correlate against. Sènica et al. [35] apply machine learning directly to this trade-off, optimizing noise radar waveforms for low probability of intercept while managing peak-to-average power ratio. Their results are in simulation, but the formulation is directly relevant to the operational constraint Thales articulated.

**Single-node geometry is a constraint, not a necessity.** Multi-frequency and multi-static geometry are useful for drone discrimination. Palamà et al. [36] demonstrate this directly. Using the NeXtRAD multistatic radar system [37] at roughly 1.3 GHz

and 10 GHz with separate transmit and receive nodes, they show that the blade-to-fuselage return ratio for drones varies between  $-11$  and  $-3.7$  dB depending on viewing geometry and frequency, while bird returns cluster differently. A monostatic, single-frequency radar gets one sample from that distribution. A multi-frequency multistatic system gets many. We focus on one radar node for now, but daydream about multistatic, in particular what RF over fiber paired with disposable fiber or transmitters may enable.

## 2.4. Cognitive Radar

The term *cognitive* is used in a broad enough manner to make the published literature hard to compare [38]. Getting the definition right matters because it determines whether the handful of hardware demonstrations represent meaningful progress or familiar adaptive signal processing in new language. Practical aspects [24] include latency bottlenecks, various perception-action cycle design considerations, and how to work in dynamically-changing environment. No standard has yet been established within the radar community.

For a radar system to be cognitive, therefore, it is a fundamental necessity for the radar transmitter to learn from **continuing** interactions with the environment and **intelligently** use the information extracted by the receiver on targets under surveillance, all of this being done **on the fly** during different phases of the target-track sequence.

Simon Haykin [39]

The emphasis is added. These three words do real work as a filter. **Continuing** rejects anything “learn then deploy” shaped, where learning stops before the system is fielded. **Intelligently** rejects approaches that lean more on physics-based design than on anything that could reasonably be called intelligence. **On the fly** rejects anything whose inputs were ever at rest, ruling out offline datasets as the primary signal.

Some authors draw on Haykin’s definition to distinguish feedforward radar, fully adaptive radar, and cognitive radar proper, the last set apart by the inclusion of memory [38]. Others, like Shannon D. Blunt and Anthony F. Martone, make no such distinction, instead explicitly stating “cognitive radar (also referred to as fully adaptive radar)”[23].

The most common pattern is to define some sort of optimization process. They then reach for a method from machine learning as a surrogate for some component in an otherwise standard radar processing chain. When reinforcement learning techniques are used, they are nearly always an *offline* process, often with the goal of finding some policy and then ending learning, sometimes presenting this as desirable. This is often a Markov Decision Process (MDP) formalization, often featuring a stationary environment despite nearly all arguments for waveform diversity being based on the idea that the environment is not stationary. Many also delegate the reward function to an operator who tunes coefficients to accomplish a mission objective. This is typically framed as a benefit.

Recent work in the reinforcement learning community questions whether the MDP is even the right foundation for agents expected to operate in open-ended environments indefinitely [40]. The *big world hypothesis* [41] makes this precise: for any agent of bounded capacity, the world will always be bigger than the model. They instead lay out a *history*

*process* framework that is attractive for this problem. This requires defining a reward with no ability to inspect state.

## 2.5. Digital Pre-Distortion

Radar transmitters are designed to operate amplifiers near saturation to maximize power efficiency. This results in a few well-known artifacts, spreading of the spectrum from odd-order nonlinearity problems, spreading of Doppler ambiguity, widening of range sidelobes, and new sidelobes close to the main lobe from third-order nonlinearities. Many pieces of work constrain themselves to maintain constant modulus [15, 21] due to the complex nonlinearities of the power amplifier.

Digital pre-distortion is an established communications technique, targeting amplifiers below saturation via model-based approaches (Volterra series, memory polynomial) or small neural networks [42]. The memory polynomial, a tractable subset of the full Volterra series, models the amplifier output as:

$$y[n] = \sum_{p=1,3,5,\dots}^P \sum_{q=0}^Q c_{p,q} x[n-q] |x[n-q]|^{p-1}, \quad (2.1)$$

where  $P$  is the nonlinearity order,  $Q$  is the memory depth in samples, and  $c_{p,q} \in \mathbb{C}$  are coefficients fit to loopback measurements. Only odd orders appear because even-order terms alias outside the signal band and vanish after baseband filtering. Radar amplifiers running in saturation are a different problem. Standard DPD requires usable gain headroom above the compressed operating point. Saturated operation removes that headroom. At the limiting extreme, many input amplitudes produce the same clipped output, so no unique inverse exists. Communications DPD also optimizes for adjacent-channel power ratio and symbol fidelity, while radar performance is set by the ambiguity function. Nonlinear distortion that is invisible to spectral containment metrics can still raise range sidelobes enough to mask slow targets [19, 43]. The nonlinearity is severe enough that most radar waveform design work just avoids the issue entirely by sticking to constant-modulus signals.

## 2.6. Takeaways

Waveform designs that outperform frequency modulation on paper are not rare, designs verified against real targets on real hardware are. This matters because simulation environments fail in 2 directions. They are too easy, often excluding well characterized amplifier distortion, real clutter statistics, unexpected interference, and dozens of other known but difficult to model impacts. They are too difficult, full of uniform random, Gaussian, or Poisson distributed issues where no exploitable strategy exists. For some classes of problems, this is adequate. But with powerful learning, simulation is weak evidence that something works, and equally weak evidence that it does not.

To our knowledge, no published cognitive radar demonstrates genuine *continual* adaptation on hardware without episode resets, hand-engineered reward functions, and a narrow degree of freedom on what parameters the radar may adapt.

The field has mostly responded to the drone detection problem with better signal processing inside the existing pipeline rather than questioning the pipeline.

**Table 2.1:** Waveform diversity work surveyed in this chapter. Simulation-only above the rule, hardware-demonstrated below.

Ref	Waveform	Platform	Scenario	Key Result
[44]	RFM	Simulation	Vehicle classification in low SNR	RFM enables aspect-angle classification in simulation
[31]	Multi-waveform set	Paper study	Sidelobe level, resolution, spectral containment	FTE and PRO-FM dominate spectral efficiency comparisons
[15]	Phase-only constant modulus	Simulation	Range-Doppler shaping against known clutter	Non-convex quartic program places sidelobes on modeled clutter map
[45]	Costas and Sudoku codes	Simulation, 8–12 GHz MIMO	Moving targets	Formal characterization of Doppler tolerance for polyphase codes
[26]	Frank codes	Simulation (GAN)	Ambiguity function synthesis	Generative model reproduces Frank-like ambiguity function statistics
[46]	RFM	Ettus B200mini + Nvidia Jetson	Traffic at 0.80 km	Real-time matched-filter pulse compression on embedded GPU
[33]	Modified OFDM	50 MHz at 3.45 GHz	Traffic at 1 km	Nonrepeating waveforms with FMCW-class range-Doppler processing
[47]	Phase-modulated CW	Automotive wideband	2 corner reflectors, anechoic chamber	Mutual interference reduced between adjacent automotive radars
[48]	Pseudo-random FM	3.55 GHz, Tektronix AWG	Traffic at 1.05–1.25 km	Dual-pol noise-FM suppresses cross-pol sidelobes over the air
[21]	L-BFGS optimized RFM	3.35 GHz carrier	Multiple moving vehicles	Optimized RFM encodes arbitrary range-Doppler signatures over the air
[22]	Sense-and-notch RFM	SDR	Moving target, hopping interferer	Hopping interferer tracked and notched in real time on hardware
[49]	Alltop waveforms	PARSAX (3.315 GHz)	Stationary chimney	Alltop ambiguity structure validated on full-power radar
[17]	Spectrally notched linear chirp	PARSAX (3.315 GHz)	Chimney, spectrum coexistence	Notched waveform coexists with incumbent, amplifier distortion model validated

# 3

## Gaps

The literature reviewed in chapter 2 has a consistent shape. Waveform design happens on paper or in simulation, with the occasional over-the-air result on a software-defined radio, nearly always offline. The 3 experiments outlined in chapter 4 each address a gap discussed in this chapter, with the full mapping in section 3.6. The same hardware opens directions that are too speculative to commit to here, collected as side bets in section 3.5.

### 3.1. Stuck in Simulation

Waveform designs are evaluated in simulation because real-hardware experiments are expensive and slow. The simulation is then made more elaborate to justify the result, and yet the elaborate simulation is not the real hardware. A large part of the reason so little of this work reaches hardware is that it requires an operational radar. One response is conceptually simple, if not easy in practice, to close that gap. Run on real hardware and learn from real returns where possible given the constraints of the specific research project.

The Delft University of Technology radar PARSAX<sup>1</sup> is software-defined, dual-polarimetric, with a 14-bit AWG at 100 W per channel and the capture chain colocated with a GPU<sup>2</sup> on the same machine. This combination at a university is the primary justification for running these experiments here rather than elsewhere. Ettus software-defined radio hardware lacks the transmit power to illuminate real targets at surveillance ranges [46], though recent work shows that pairing an Ettus radio with an Nvidia Jetson is capable of matched-filter pulse compression on random frequency-modulated waveforms in real time [46].

### 3.2. Pre-Distortion Mismatch

Digital pre-distortion has a mature literature in communications, with dozens of papers on Volterra-series models, memory polynomials, and learned variants trained end-to-end [42, 50]. The implicit operating assumption throughout is that the amplifier is running near its linear region, where distortion is an unwanted perturbation and the goal is to cancel

---

<sup>1</sup>In this work, only the 3.315 GHz channel is used.

<sup>2</sup>GPU here denotes any high-throughput parallel compute device; the distinction between NPU, TPU, and GPU branding is irrelevant to the work.

enough of it that spectral mask requirements and error-vector magnitude targets are met. Radar amplifiers do not necessarily share this assumption. The distortion is not a small perturbation on a nearly-linear system. It is a large, operating-point-dependent deformation of the whole output [43].

The performance metric is also different, and the difference matters more than the operating point. Communications pre-distortion is validated against spectral mask compliance and error-vector magnitude, asking whether the signal stays inside its allocated band and whether the symbol constellation looks clean. Radar waveforms may have no constellation. What matters is the ambiguity function, the 2-dimensional shape of the range-Doppler response, including the sidelobe structure that determines clutter rejection and the main-lobe width that determines resolution. Amplifier nonlinearity reshapes the ambiguity function in ways that spectral mask compliance does not capture [19]. A pre-distortion scheme that satisfies a communications metric does not necessarily help the range sidelobe structure. This means the communications digital pre-distortion literature offers no validated method for radar, leaving an interesting gap to fill.

### 3.3. Continual Learning

Published cognitive radar demonstrations typically train a policy, freeze it, and deploy it. This amounts, in the end, to a fixed waveform with extra steps. The genuinely hard problem is a radar that keeps learning after deployment, against an adversarial operator who knows it is learning, in a non-stationary environment, without the opportunity to reset or retrain. Elelimy et al. [40] and Javed and Sutton [41] frame why this is hard and point toward history-process formulations as a credible attack. Demonstrations in hardware are scarce in the literature.

In the reinforcement learning community, these big world problems are computationally demanding and under-benchmarked. Hardware-in-the-loop radar is an attractive test bed: waveform changes alter the statistics of the collected data almost immediately, producing a genuinely non-stationary stream that offline replay cannot replicate.

This gap motivates the architecture but is not closed by this thesis. Closing the hardware loop is the precondition, and the Sense-and-Notch online-updating extension is the first concrete step in that direction.

### 3.4. Trustworthy Labels

Any supervised approach requires labels in sufficient quantity and quality. For drone tracking the practical approach is a GPS over Long Range, a low-throughput sub-GHz radio protocol (LoRa) setup. LoRa is a radio protocol for long-range low-throughput communication on cheap hardware. Fitting a position string into a LoRa payload requires truncating the decimal places in the coordinate, which sets accuracy, and reducing the update rate to respect LoRa's duty-cycle limit, which sets the rate. The tracker gives position at 1 Hz with 3 m horizontal accuracy. This confirms which range bin the drone is in. It is not accurate enough to supervise learning of fine-grained radar-to-position mappings or to synchronize multi-radar observations to a common reference. The gap is what motivates

a self-supervised objective built from the radar's own returns, rather than from external labels.

## 3.5. Speculative Directions

The hardware assembled for these experiments opens directions that are interesting but too far from the core question to commit to in this proposal. They are recorded here as side bets, distinct from the gaps the experiments close.

### 3.5.1. Adapting Only the Transmitter

Often the literature only optimizes the transmitter. The receiver is fixed: matched filter, constant-false-alarm-rate threshold, standard pipeline. Jointly adapting both requires the transmitter and receiver to negotiate, and finding the combination that works best is not decomposable into 2 separate optimization problems. This is consequential. A reinforcement learning agent trained over the joint action space of waveform and processing mode, including real-time full-polarimetric processing, in a closed hardware loop is the natural approach, and the hardware to attempt it exists. The signal processing is understood in principle, the bottleneck is compute throughput at the required pulse repetition intervals, and a GPU today has the memory bandwidth to handle it. The bootstrapping is the real difficulty to address this gap effectively. The transmitter and receiver objectives couple in ways that resist easy decomposition, and no clear attack exists yet.

### 3.5.2. Radio-Frequency Tokenization

Transformers [51] have displaced prior state of the art in language, vision, robotics, and games, not by encoding domain expertise but by scaling general-purpose pattern matching against large token sequences. Radar problems have the right shape for this. Returns mix too many interacting effects to model cleanly from first principles, and a common response is to declare the residual noise and whiten it. The current generation of attempts, with RF-GPT as an example [52], takes the obvious path of rendering a spectrogram and passing it to a vision model. This treats the transformer as a human operator who has learned to read range-Doppler maps by eye. A transformer is not that. It predicts sequences of tokens, and it does this well when the tokenization is helpful, so the most informative path is unlikely to run through red-green-blue pixels. Modern reinforcement learning world models do not ask a transformer to predict raw pixels. They compress the environment into a latent space, tokenise there, and let the transformer operate entirely in that space [53]. The joint transmitter-receiver adaptation problem of section 3.5.1 has exactly that structure, where a history of transmitted waveforms and received returns determines the next waveform. A good radar tokenizer turns that into a standard sequence-modelling problem. Radar is, despite the community's frequent claims to the contrary, not data-limited. Radars are data-generating machines.

### 3.5.3. Motion-Capture Ground Truth and Multistatic Localization

The ground truth necessity of section 3.4 has an indoor alternative. The Cyber Zoo of Delft University of Technology is an indoor flight arena instrumented with an OptiTrack motion capture system, which tracks retro-reflective markers to sub-millimetre accuracy at over

100 Hz.<sup>3</sup> Position and attitude ground truth there is 3–4 orders of magnitude better than GPS outdoors. The catch is range. The Cyber Zoo is enclosed, so targets sit at 2–20 m, and the radar cross section at those ranges is not representative of a 400 m engagement. It is a place to train and validate learned tracking before moving outdoors. Delft University of Technology also has a benchtop AWG independent of PARSAX, which combined with a corner reflector on a track is an even simpler controlled setup and permits much longer continuous measurement sessions than outdoor field campaigns allow.

A single radar resolves range and radial Doppler. 2 radars at different angles add a second range measurement and a second Doppler projection. 3 radars with wide angular separation can triangulate a drone’s 3-dimensional position from range measurements alone, without any phase synchronization between nodes. The multistatic radar literature is substantial, and fusion algorithms exist for combining range-only measurements from multiple non-coherent nodes [54]. What does not exist is a model trained end-to-end from raw samples at multiple non-coherent nodes to a fused position estimate, supervised against high-accuracy ground truth.

The sensors are the missing piece, and we are in early discussion with Dopplium<sup>4</sup> about deploying their radars around the Cyber Zoo perimeter for this purpose.

The natural experiment is then to place those radars around the Cyber Zoo perimeter and observe drones during other groups’ flights, with OptiTrack as ground truth. A model trained end-to-end from raw samples to position could discover features that a geometric model ignores, including multipath and near-field structure. This is not one of the 3 proposed experiments, it is the step that becomes feasible once the outdoor pipeline validates. Prior work establishes that multiple non-coherent radar views improve drone classification [54] and that distributed fusion improves target location and velocity estimates [55]. Extending that to 3-dimensional localization with end-to-end training and sub-millimetre ground truth remains open.

## 3.6. Scope

Three gaps are addressed directly by the lines of work in chapter 4. All three run on PARSAX against real returns, which closes the simulation-to-hardware gap of section 3.1 to the extent that a single facility allows. The pre-distortion mismatch of section 3.2 is the central problem for the learned DPD experiment. The ground truth necessity of section 3.4 motivates Target-Matched Illumination, which avoids the label problem by constructing a self-supervised objective from the radar’s own returns. The continual learning gap of section 3.3 is not closed, but it is the explicit motivation for the Sense-and-Notch online-updating extension, the first hardware step in that direction.

Joint transmitter-receiver adaptation of section 3.5.1 is not attempted here. It is recorded as a speculative direction because the coupling between transmitter and receiver objectives resists easy decomposition and no clear attack exists yet.

---

<sup>3</sup>The Cyber Zoo <https://tudelft.nl/en/ae/organisation/departments/control-and-operations/facilities/drone-facilities/cyberzoo> is an indoor unmanned vehicle research facility in the Aerospace Engineering faculty at Delft University of Technology, where regular drone operations take place.

<sup>4</sup><https://www.dopplium.com/>

---

Other gaps are real but not attempted here. Multiple simultaneous targets in the same range-Doppler cell, against non-Gaussian clutter with unknown statistics that change between pulses, remain open in recent detection theory [56] and are among the scenarios Thales describes as operationally important. The tools exist, with adaptive Constant False Alarm Rate (CFAR) variants, sparse recovery, and Bayesian clutter modelling among them, but no convincing hardware result does yet. True multistatic coherent processing requires tight timing and phase synchronization across nodes, infrastructure that exists in only a handful of well-funded programs. The non-coherent localization bet of section 3.5 reaches part of the way there without that infrastructure, but it too is left for after the single-node pipeline validates.

# 4

## Feasibility Study

This chapter develops three lines of work that address the gaps in chapter 3. The overarching aim is to move the waveform from a parameter specified by hand toward a quantity learned from data and from the hardware itself. The three lines of work move toward that aim, from the most heavily guided to the least. Learned Digital Pre-Distortion (section 4.2) is handed a target waveform and must invert the impacts of the hardware to reproduce it. Sense and Notch (section 4.3) removes the target waveform and decides on the transmit waveform based on the return. Target-Matched Illumination (section 4.4) goes further still, discovering from raw returns which waveform properties best reveal a target.

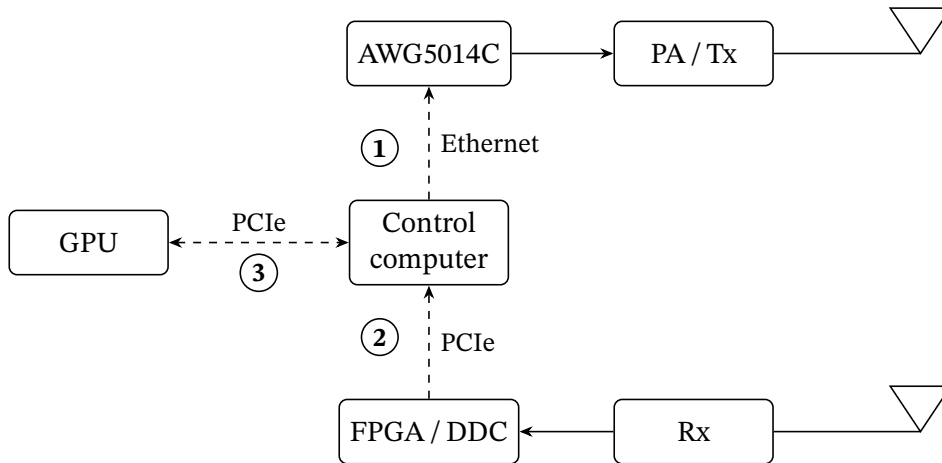
Each line of work is presented in the same structure. First the problem and its relation to the existing literature, then the general framework proposed to solve it and how that framework is validated in simulation, and then how it is implemented on the PARSAX hardware. The experiments within each line of work prove and validate the framework.

The common theme is that the radar holds the history of what it has transmitted and received, chooses the next waveform, observes the return, and updates. The lines of work differ in how much of that choice is specified by hand. Pre-distortion is handed the exact waveform and learns to realize it through the hardware, Sense and Notch is told what to remove and learns where, and Target-Matched Illumination is told only that a target should be revealed and learns how. The loop is stated formally as the illumination objective in equation (4.3), the most general problem in this chapter and the one the earlier lines of work build the machinery to attack.

The three lines of work share a prerequisite, the closed-loop infrastructure described in section 4.1. Pre-distortion validates that the hardware round-trip produces the intended waveform shape, something which Sense and Notch then relies on. Sense and Notch validates the proposed closed-loop reinforcement learning infrastructure, which Target-Matched Illumination reuses with a more general reward signal.

### 4.1. Close the Loop

**Hypothesis.** The PARSAX system can upload arbitrary waveforms from a GPU-attached host, receive returns, and iterate on a new waveform within 1 second, end-to-end.



**Figure 4.1:** Current PARSAX hardware interconnect. Dashed callouts (1), (2), and (3) mark the integration gaps.

**Contribution.** The individual components in PARSAX exist but the integrated loop does not. This is not a research contribution by itself, but without it the following questions remain unanswerable, making this an essential research tool. The hardware round-trip components are already individually verified section 4.2.4. The remaining “close-the-loop” work is the FPGA firmware modification that removes the pulse-length constraint, and demonstrating processing returns on the GPU.

Three connections separate the current state from the closed loop all three lines of work require, marked as dashed callouts in figure 4.1.

1. The AWG is reachable over Ethernet for waveform upload, but the interface is designed for control traffic. Standard python libraries do not interact with it, but it is simple and well documented control sequences.
2. The FPGA capture card connects to the control computer via Peripheral Component Interconnect Express (PCIe), but the current path is based around capturing manual “snapshots” to disk and downstream processing reads from disk on a separate machine, placing the loop update rate at (relatively slow) human-operator timescales.
3. The GPU sits on the same PCIe bus but all matched filtering and digital down-conversion run on the FPGA. Reformatting the firmware pipeline for direct GPU access should be possible, but is not directly supported.

PARSAX is an excellent platform for experimentation. We ran an initial experiment (chapter 5) with PARSAX to evaluate uploading custom waveforms, flying and illuminating a drone, and processing returns. This experiment revealed two issues with the system as it exists, which must be addressed in order to run any of the proposed experiments. We have since run another experiment (section 5.3), where we uploaded waveforms over the Ethernet interface and captured returns with the newer FPGA capture card. Initial results confirm that waveform parameters were well selected, but surfaced other hardware issues which are not within the scope of this thesis.

The FPGA currently runs waveform-specific processing that must be rebuilt for each pulse length, the engineering details and work already done to address this are in chapter 5.

If the FPGA firmware cannot be changed on the required timeline, there is a Plan B: fix the pulse length to  $2^{16}$  samples at 500 MP/s (131  $\mu$ s) and work entirely within that constraint for the first 2 experiments, relying on offline replay of pre-recorded captures for initial model development before the real-time loop closes. This limits waveform diversity, but does not block the core research question.

## 4.2. Learned Digital Pre-Distortion

**Hypothesis.** A neural network conditioned on a desired waveform can learn to produce the pre-distorted “command waveform” that, after passing through the power amplifier, recovers the desired waveform, within specifications that a fixed Volterra-series calibration on this hardware may struggle to hold as the amplifier heats [19, 20].

### 4.2.1. Problem and Prior Work

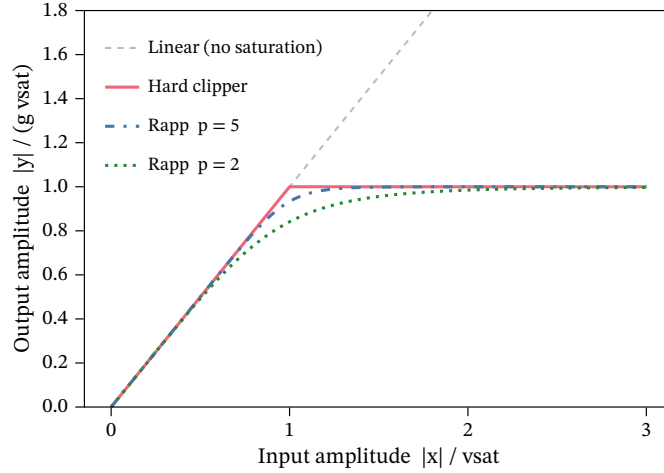
Non-linearities at saturation are unavoidable, unique to each physical unit, and change as the hardware gets hot or switches power modes. The consequences appear in the transmitted waveform: spectral sidelobes spread, notches fill back in, and whatever property was designed into the waveform gets smeared by the hardware.

Digital pre-distortion is well studied in communications, where the goal is to linearize high-peak-to-average-power-ratio signals [42]. Radar amplifiers run near or at saturation by design, and the waveforms of interest are often constant-modulus, which eliminates peak-to-average constraints as a design handle. Aubry et al. [19] jointly optimize the transmit waveform and Volterra series pre-distortion coefficients on PARSAX, but the calibration is a snapshot, it must be reset as hardware changes state. The approach proposed here learns instead continuously from loopback measurements, requires no analytic amplifier model, and is intended to track a non-stationary hardware environment rather than assume a fixed operating point.

### 4.2.2. Framework


Our proposed approach treats pre-distortion as conditional waveform generation. Given a desired waveform, a neural network is trained to return the pre-distorted version that, after passing through the amplifier, comes out closer to the desired waveform. This is, intentionally so, the most heavily guided line of work in this report. The ambition for the later lines of work is to skip the explicit waveform specification entirely, going directly from a detection objective to a hardware-ready waveform. Pre-distortion establishes that the network can learn the hardware’s forward map and invert it. This is a necessary precondition for everything that follows.

With hardware, we treat the amplifier as a black box. When simulating, we initially took advantage of being able to back-propagate gradients through it, which constrains the choice of amplifier model for simulation. The soft limiter, the standard in waveform design literature and the model whose clipping constraint appears in most waveform optimization work, is not differentiable at its saturation point. Gradients vanish as the amplifier is driven into saturation. The Rapp smooth saturation model [57] avoids this. Its output  $y(t) = g x(t) \left(1 + (|x(t)|/v_{\text{sat}})^{2p}\right)^{-1/(2p)}$  is smooth everywhere, linear at small input, and



**Figure 4.2:** Amplitude transfer curves for the hard clipper and the Rapp model at two knee sharpness values, normalized so that  $v_{\text{sat}} = 1$ . At  $p = 2$  the compression begins well below saturation, which is representative of the PARSAX Gallium nitride (GaN) solid-state amplifier. At  $p = 5$  the curve is nearly flat until the knee, then falls sharply, approaching the hard clipper as  $p \rightarrow \infty$ . Gradients of the hard clipper are zero at saturation. The Rapp model is smooth everywhere, which is essential for gradient-based training through the simulated amplifier.

saturates toward  $\pm g v_{\text{sat}}$  as the input grows. Figure 4.2 shows the transfer curve for  $p = 2$ ,  $p = 5$ , and the hard clipper resulting from  $p \rightarrow \infty$ .



**Figure 4.3:** Signal chain for learned pre-distortion. IQWaveNet maps the desired baseband waveform  $x_{\text{ref}}$  to a pre-distorted command  $x_{\text{cmd}}$ , which the AWG upsamples  $8\times$  to a real intermediate-frequency signal at 500 MS/s, amplifies through the Power Amplifier, and the Digital Down-Conversion decimates back to 62.5 MS/s. The dashed bypass shows the uncompensated path to be compared with the proposed approach;  $\text{PA}(\cdot)$  denotes the full chain output used throughout this work. The digitized output of this chain is  $\text{PA}(x_{\text{ref}})$ , which we use directly as  $x_{\text{obs}}$  in the training loop.

### 4.2.3. Simulation

The work accepted to EuRAD 2026 [58] (reproduced in chapter A) establishes the approach in simulation. The key results are summarized in section 4.2.3.

#### 1D U-Net

The submitted architecture, shown in figure 4.3, takes waveform  $x_{\text{ref}}$  as input and outputs waveform  $x_{\text{cmd}}$ , both real-valued sample arrays at 500 MP/s quantized to 14-bit integers before upload to the AWG. The backbone is a 1D U-Net [59] (a convolutional encoder-decoder with skip connections) conditioned on  $x_{\text{ref}}$  at each scale, embedded in a discrete diffusion process to generate the pre-distorted waveform over multiple denoising steps.

### Simulation loop

The simulated hardware path wraps the Rapp model with a loopback attenuation of roughly 15 dB, a fixed sample delay representing cable length, and additive white Gaussian noise calibrated to a 40 dB SNR at the digitizer. In figure 4.3 this is hidden away in the DDC block. The training loop proceeds as follows.

1. Pre-generate a library of reference waveforms: LFM chirps at a range of bandwidths and sweep times, and notched variants.
2. Draw  $x_{\text{ref}}$  from the library and pass it through the current network to produce  $x_{\text{cmd}}$ .
3. Pass  $x_{\text{cmd}}$  through the simulated Rapp path to obtain  $x_{\text{obs}}$ . An early tap ( $x_{\text{awg}}$ ) is logged for diagnostics. In hardware this step is replaced by the real AWG-to-ADC chain with no other change to the code.
4. Compute the combined loss:

$$L = c_t \left( \frac{1}{N} \sum_i |x_{\text{ref}}[i] - G x_{\text{obs}}[i]|^2 \right) + c_f \|\text{PSD}(x_{\text{ref}}) - \text{PSD}(x_{\text{obs}})\|_2^2 \quad (4.1)$$

The scalar  $G$  accounts for loopback coupling attenuation and is estimated from the data. Aubry et al. [19] perform a least-squares time-domain alignment step before comparing waveforms. The spectral term is computed over the signal bandwidth window. For notched waveforms (section 4.3), the primary failure mode is a notch filling back in under amplifier saturation, and a large  $\|\text{PSD}(x_{\text{ref}}) - \text{PSD}(x_{\text{obs}})\|_2^2$  catches it directly.

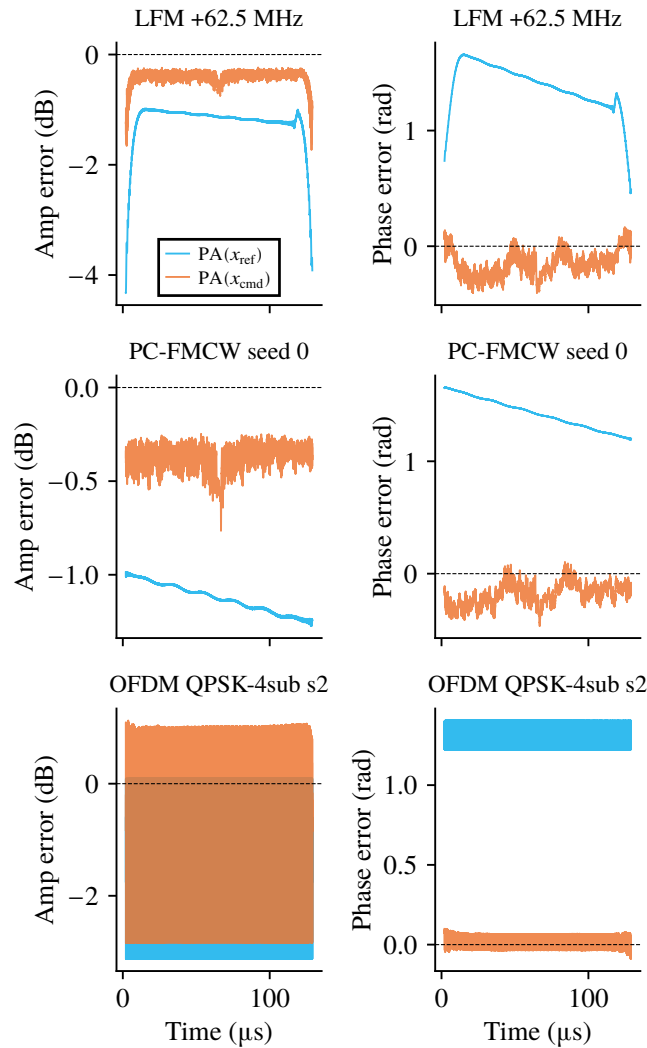
5. Run multiple models in parallel with different random initializations and loss weightings ( $c_t, c_f$ ), promoting the best-performing one each update cycle following the population-based training scheme of Jaderberg et al. [60].
6. Repeat from step 2 until convergence.

### Results

A single network with no waveform-specific parameters reduces root-mean-square distortion between the reference and post-amplification waveform by 67–75% for LFM chirps, 72% for smoothed phase-coded waveforms, and 11–61% for OFDM waveforms at the amplifier saturation knee, with error profiles shown in figure 4.4. The wide spread on OFDM reflects peak-to-average power ratio: high-subcarrier OFDM compresses more and sees larger absolute improvement.

The more exciting result is generalization. We group waveforms into broad “families” of being either basic LFM waveforms, phase coded LFM waveforms, or OFDM waveforms. Training on the 2 LFM waveform families and holding the third out entirely still achieves 51% improvement on the unseen, and fairly different, OFDM family. This suggests the network learns the amplifier’s response rather than a per-waveform correction. That is the property needed for a radar that encounters waveform types not present during calibration. An important caveat is that smoothed phase-coded waveforms and OFDM share a high peak-to-average power ratio, so the 51% transfer may partly reflect that shared structure rather than true cross-family generalization. Waveform families that differ more substantially

in their envelope statistics would be a harder test. It is relatively easy to feed waveforms through the simulator and measure this same loss function, but more interesting are waveforms with radar specific applications where we can show improved performance with pre-distortion.



**Figure 4.4:** Amplitude and phase error across the pulse duration for one example per waveform family, with and without pre-distortion. The uncorrected baseline (blue) shows a systematic amplitude ramp and phase drift that accumulate over the pulse. The network output (orange) is flat and small throughout. The error pattern is a signature of thermal memory in the amplifier; the network anticipates it rather than reacting to it, requiring the receptive field to span the full waveform duration.

Three extensions stay entirely in simulation. First, a harder cross-family test, holding out a waveform family whose envelope statistics differ substantially from the training families rather than ones that share a high peak-to-average power ratio, to separate true cross-family transfer from shared structure. Second, a sweep over amplifier operating points, knee sharpness, gain, and thermal state, to measure how far a single network generalizes across the non-stationarity it is meant to track. Third, the same study run on the original diffusion architecture, so that the generalization of the slower network is a fair baseline against the faster IQWaveNet rather than an uncontrolled change. These extensions need no hardware and can proceed in parallel with the closed-loop work below.

### 4.2.4. Hardware Implementation

Moving from the submitted simulation to real hardware requires two departures from the EuRAD 2026 setup.

#### Training without a differentiable amplifier

In simulation, the Rapp model is smooth everywhere, including at saturation, so gradients flow through it without issue. A physical power amplifier offers no gradient. We propose to use iterative learning control, framed as supervised regression, to sidestep the requirement and address the problem. This will be evaluated both in simulation, initial results shown in figure 4.6 are encouraging, and most importantly in hardware. Each iteration proceeds as follows.

1. Pass  $x_{\text{ref}}$  through the network to produce  $x_{\text{cmd}}$ .
2. Transmit  $x_{\text{cmd}}$  via the AWG. The FPGA capture returns the observed output  $x_{\text{obs}}$ . The amplifier is opaque to the optimizer.
3. Compute the residual  $e = x_{\text{obs}} - x_{\text{ref}}$ , the error the network failed to compensate.
4. Form the corrective label  $x_{\text{target}} = \text{sg}(x_{\text{cmd}}) - Ge$ , where  $\text{sg}(\cdot)$  stops gradient flow through the current command.  $G$  is a scalar learned from data to account for simulated attenuation.
5. Take a gradient step minimizing  $\|f_{\mathbf{w}}(x_{\text{ref}}) - x_{\text{target}}\|^2$ , treating  $x_{\text{target}}$  as a fixed label. Repeat from step 1.

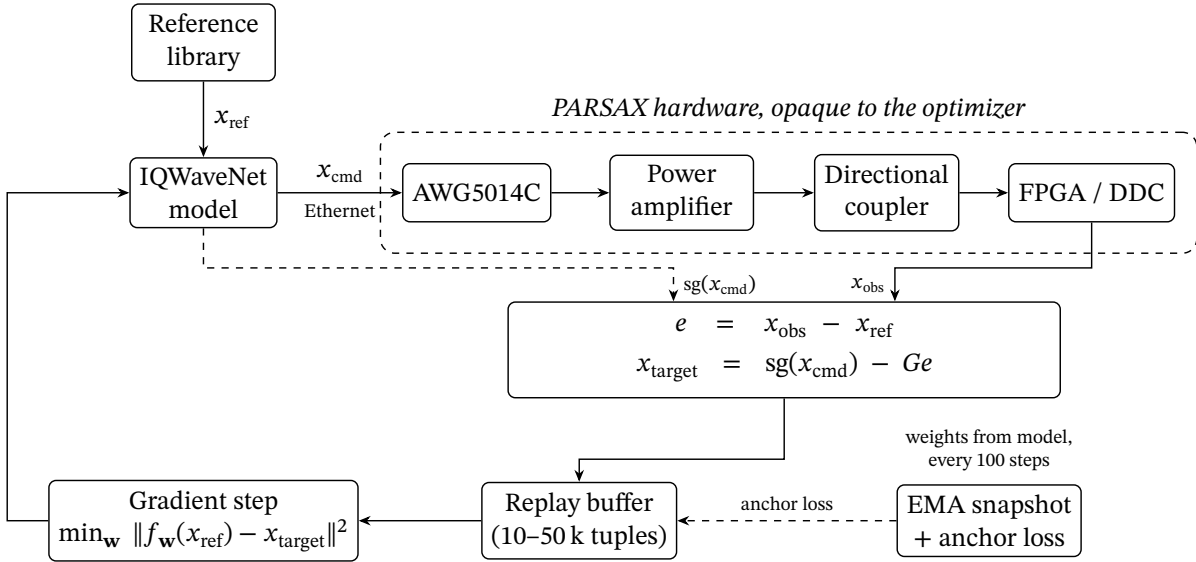
The scheme stays on-policy: each label derives from the current network running on the current hardware, so the training signal tracks the amplifier as it heats and drifts.

The gradient step for iteration  $k$  is:

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \alpha \nabla_{\mathbf{w}} \left\| f_{\mathbf{w}_k}(x_{\text{ref}}) - x_{\text{target}} \right\|^2, \quad (4.2)$$

where  $\alpha$  is the learning rate and  $x_{\text{target}} = \text{sg}(x_{\text{cmd}}^{(k)}) - Ge^{(k)}$  is treated as a fixed label with no gradient flowing through it.

Figure 4.5 shows the full loop. The top row is the forward pass, left to right:  $x_{\text{ref}}$  feeds the neural network,  $x_{\text{cmd}}$  travels over Ethernet to the AWG, passes through the power amplifier and coupler, and the FPGA returns  $x_{\text{obs}}$  via shared memory. The dashed border marks the physical hardware, which is opaque to the optimizer. In the lower half, the iterative learning control block receives  $x_{\text{obs}}$  on a solid arrow and  $\text{sg}(x_{\text{cmd}})$  on a dashed arrow; the dashed path is the stop-gradient, the point at which the amplifier exits the computation graph. The corrective label feeds the replay buffer, which holds the most recent 10–50k  $(x_{\text{ref}}, x_{\text{target}})$  tuples. Training on a rolling window rather than a single observation per step smooths the gradient signal and limits how much a single anomalous hardware return can shift the weights. The gradient step block draws random batches from the buffer and updates the network by minimizing the supervised regression objective. The exponential moving average (EMA) snapshot holds a slowly drifting copy of the network weights, refreshed every 100 steps. Its anchor loss penalizes the current network for diverging from the EMA snapshot's predictions on the same inputs, limiting how fast the weights can change in response to a sudden hardware transient without discarding a solution that was working.



**Figure 4.5:** Hardware-in-the-loop iterative learning control training loop for learned pre-distortion. Solid arrows carry the forward pass; dashed arrows mark stop-gradient paths. The dashed border encloses the physical hardware, which is opaque to the optimizer.

### Neural network architecture

The EuRAD 2026 submission initially explored a discrete diffusion process to generate pre-distorted waveforms. That approach requires sampling over many denoising steps, which is slow for a hardware loop that must complete within 1 second per iteration. The intended updated network is IQWaveNet: a 1D residual WaveNet with 9 dilated convolutional blocks (dilations 1, 2, 4, ..., 256) interleaved with a single bottleneck self-attention layer at the midpoint. The network maps the complex-valued IQ chirp  $x_{\text{ref}} \in \mathbb{R}^{2 \times 8192}$  directly to the command waveform  $x_{\text{cmd}}$ , initialized as the identity so training starts from pass-through behavior. A single forward pass at inference takes under 2 ms on the RTX 5090, well within the 1-second round-trip budget.

### Hardware interface

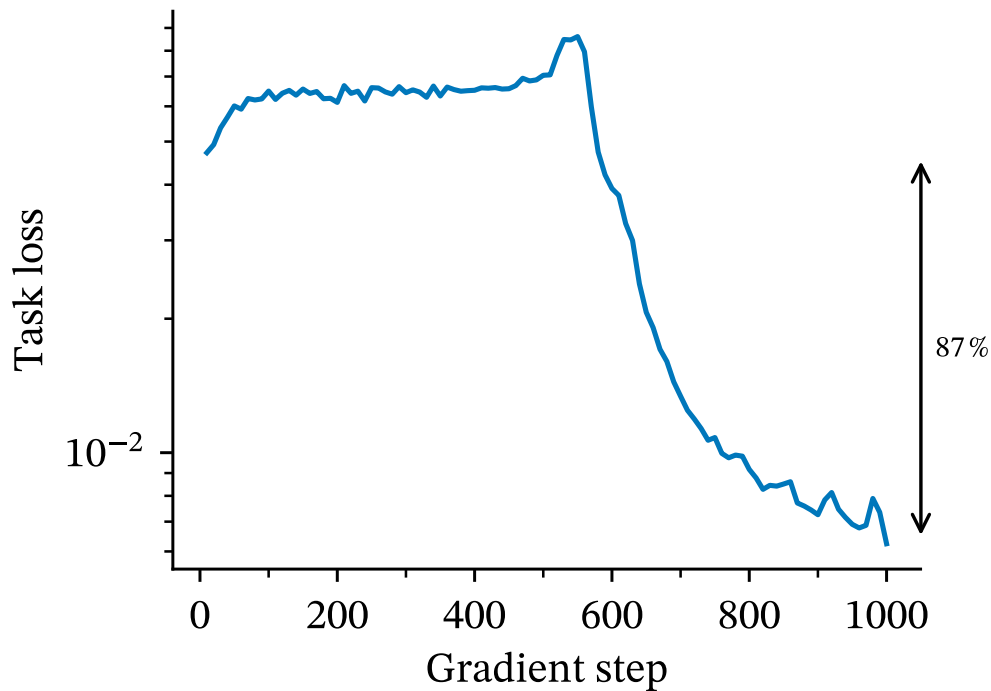
Thomas Epailly (MS<sup>3</sup> Laboratory, Delft University of Technology) has written 2 software libraries that provide the hardware interface [61]. `awg5014_python` is a TCP client for the AWG5014C that accepts a 14-bit integer sample array and handles binary framing, waveform naming, and channel routing. `ptshm-python` is a C extension that reads PARSAX FPGA capture data from a shared memory region maintained by `parsaxtool`, returning one ( $N_{\text{DDC}}, 2$ ) int32 snapshot per captured pulse. `awg5014_python` is already in field use, as it drove waveform upload for the second drone measurement campaign at Poeldijk (discussed later in section 5.3), replacing the manual front-panel process from the first experiment. `ptshm-python` is being integrated into the hardware-in-the-loop training loop now.

### Closed loop verification

The full hardware round-trip works end-to-end. A waveform generated in JAX<sup>1</sup> on the PARSAX machine’s RTX 3080 and quantized to 14-bit integers plays on the AWG, and the

<sup>1</sup><https://github.com/jax-ml/jax>, a Python library for GPU-accelerated array computation.

corresponding FPGA capture lands in GPU-host memory via ptshm. The training loop runs in simulation with the same code path, and figure 4.6 shows that the loss falls 87% over 1 kSteps despite the optimizer never accessing a gradient through the amplifier. The loss rises slightly for the first 200 steps before falling; the replay buffer is still sparse at that point, so gradient steps train on nearly the same examples repeatedly and the corrective labels are noisy until enough diversity accumulates. Moving to hardware should, in principle, amount to substituting the simulated amplifier path for the real AWG-to-capture connection, though we do not expect the transition to be quite that clean.



**Figure 4.6:** IQWaveNet task loss under iterative learning control, simulated Rapp amplifier, with no gradient access to the amplifier. Loss falls 87% over 1 kSteps.

One open question from simulation is whether the Rapp model is causing enough distortion for pre-distortion to matter in a downstream radar pipeline. The simulated Gallium nitride (GaN) amplifier may not produce sufficient distortion on OFDM waveforms to motivate pre-distortion. Constellation plots from the EuRAD Student Challenge<sup>2</sup> GaN challenge are consistent with this: the simulated nonlinearity appears mild on OFDM. A real amplifier running near saturation is expected to resolve this ambiguity by producing more significant distortions and effects to be compensated, and the hardware experiment is the test.

### Proposed Experiment

The hardware milestone is latency. We will evaluate how well the IQWaveNet and the iterative learning control scheme transfer from simulation to actual amplifier data. The round-trip time from the ADC capture completing to a new  $x_{\text{cmd}}$  playing on the AWG must fit within 1 second. That budget covers GPU inference (a single forward pass through

<sup>2</sup><https://www.eumw.eu/general-information/student/student-design/> which on 2026-06-03 hosts a page titled “The GaN challenge: mastering OFDM radar under real PA nonlinearities”

IQWaveNet, under 2 ms), 14-bit quantization, Ethernet upload (about 1 millisecond for a  $2^{16}$ -sample waveform), and PCIe transfer of the digitized return from the FPGA to the GPU. The PCIe transfer is the one piece that the first deliverable described in section 4.1 must resolve before this benchmark is meaningful. Once the FPGA pipeline is offloaded, it should become a predictable and fast step.

The principal open question is whether the iterative learning control scheme tracks amplifier thermal drift fast enough on actual hardware. If it cannot track drift at the required rate, the fallback is periodic offline recalibration at longer timescales, still more robust than the single-snapshot calibration of prior work [19] but without the continuous tracking property. Simulation cannot answer this; the hardware experiment will.

## 4.3. Sense and Notch

**Hypothesis.** An agent observing only the raw return from the radar can decide how to transmit to suppress interference that masks a target, without prior knowledge of the interference. The first experiment narrows this to a single, reproducible interference type. A reinforcement learning agent tracks a frequency-hopping interferer and selects a pre-notched waveform that suppresses it, without being told the interferer’s hopping sequence in advance.

### 4.3.1. Problem and Prior Work

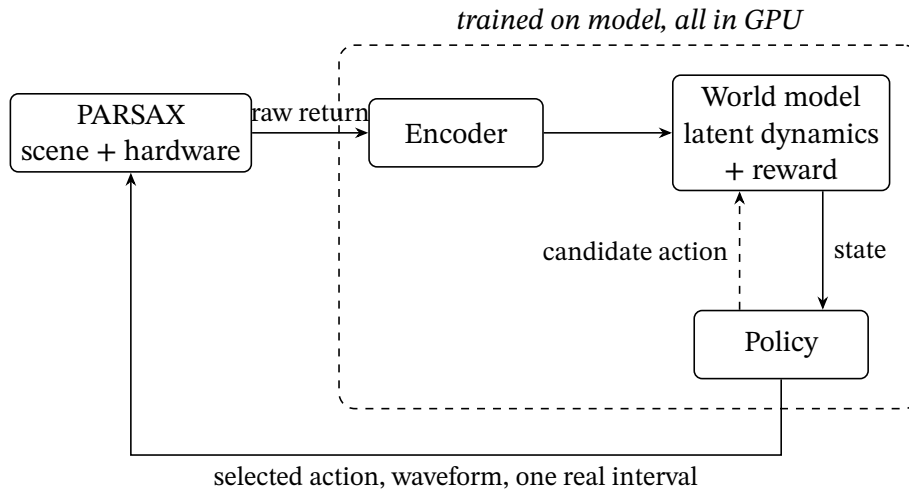
The pre-distortion line of work hands the neural network a target waveform and asks it to invert the hardware. The next piece of work involves making the specific waveform no longer specified, it is decided from the return.

Spectral notching in PARSAX has been demonstrated as an offline waveform design step [17], where the notch is placed before the experiment and stays fixed. The sense-and-notch literature handles dynamic interferers [22, 23], but those systems rely on explicit spectral sensing to locate the interferer before placing the notch. The key difference here is that the agent observes only the received signal, infers the interferer’s current frequency from that, and acts within the pulse repetition interval. There is no prior communication with the interferer and no probing waveform. The agent must infer and respond from the return it was already going to collect.

### 4.3.2. Framework

We propose a world-model agent in the spirit of Dreamer-4 [62]. A generative model learns, from logged returns, how the scene and hardware turn a transmitted waveform into a received one. The policy trains inside that learned model, exploring candidate waveforms against predicted returns. Real transmissions update the model where its predictions diverge from the observed return. Figure 4.7 shows the envisaged arrangement.

The obstacle to implement this approach is the size of the action space. A waveform is thousands of complex samples and the return is comparably large, so an unconstrained map from one to the other cannot be learned from the handful of cycles real hardware affords.



**Figure 4.7:** World-model training loop for deciding waveforms from raw returns. The encoder converts each return to a latent representation, the world model predicts the next return and reward from that state and a candidate waveform, and the policy trains on rollouts from the model. Solid arrows carry real transmissions, spent only to correct the model, the dashed arrow marks a model rollout.

The choice of Proximal Policy Optimization (PPO) as the learning algorithm sits in some tension with the critique elsewhere in this document of MDP-framed cognitive radar. PPO is an on-policy MDP method, and the Markov property it assumes does not hold in the full continual-learning problem (section 3.3). The defensible position is that PPO is a *known-good* tool for the bounded sub-problem (one coherent processing interval to the next, within a fixed interference scenario) and serves as a useful baseline even if the larger problem ultimately requires a history-conditioned formulation. Developments in the continual learning community are interesting, and should one emerge, a *known-good* algorithm from this space would be a better baseline.

A library of waveforms is pre-generated with spectral notches at various frequencies and depths, following the notch design approach already validated for PARSAX [17]. The notch library covers the operating band at a resolution finer than the interferer’s frequency grid, so that every hop position has a corresponding library entry capable of suppressing it.

### 4.3.3. Simulation

A PPO agent [63] trains offline on recorded or simulated replays of the hopping interference before the live experiment. At deployment, the frozen policy selects a notch from the library based on the current received signal, inference must complete within the pulse repetition interval, between one transmission and the next. The state representation (spectral summary, range-Doppler descriptor, or raw complex samples) is left open pending latency experiments.

This is the “learn then deploy” pattern this document criticizes elsewhere. It is the right starting point though. The question that needs answering first is whether any trained policy outperforms no notching at all on real hardware. The harder variant, where reward derives only from observables and the policy updates online without oracle information, is the first extension described below, the frozen-policy experiment is its prerequisite.

### 4.3.4. Hardware Implementation

The experiment scenario to prove this idea is PARSAX running standard range-Doppler processing against simple reference targets, a chimney, a motorway, possibly cyclists. Partway through, a software defined radio positioned at a known location near the radar activates and begins hopping across frequencies following a pre-determined sequence. The sequence is deterministic and reproducible, with fixed length, fixed hop rate, and fixed frequency grid, so the same interference scenario can be replayed exactly across experimental conditions. This pre-determined sequence is used at train time as input for a genie-informed reward signal, without any real-time communication between the interferer and the radar.

The scenario runs under three conditions, no notching as a baseline, a genie-informed *notcher* that knows the interferer's true center frequency at each step and selects the matching library entry (it has the correct frequency but not a tailor-made notch shape), and the trained PPO *notcher*. The genie condition sets an upper bound, the gap between it and the trained agent is the cost of inferring the hopping pattern from observations rather than prior knowledge. Latency of the deployed policy is reported alongside suppression performance, since spectral agility in real time is the point.

### 4.3.5. Sub-Problems and Extensions

The initial experiment is the first and most constrained piece of this line of work. The sub-problems below extend it in two directions. One is how much oracle information the training signal carries and how capable the interferer is. The other is the action space, moving from a library to total freedom.

#### Removing the genie


In the baseline design the interferer's hopping sequence is known at training time and embedded in the reward signal, the agent learns from observations, but the training signal has oracle information built in. The operationally possible variant strips this. Reward derives only from a measurable proxy (spectral power in a reference channel, or signal-to-interference ratio estimated from a range bin known to contain no target), and the agent has no advance information about the hopping pattern. This is harder, and it is the setting that matters. Whether a purely observation-driven agent can match the genie-assisted one on this hardware is the actual open question.

#### An adversarial interferer

This extension replaces the fixed-sequence interferer with an adversarial one. A jammer that can infer the radar's notch placement from the transmitted waveform will hop away from it, the fixed-sequence assumption becomes unrealistic against a sophisticated opponent. The problem becomes a pursuit game. The radar must either randomize notch placement to deny the jammer a predictable target, or predict the jammer's next frequency and notch ahead of it. Whether a learned mixed strategy outperforms a fixed pseudorandom one in this regime is open.

#### Designing the notch on the fly

The library in the initial experiment is precomputed, and every entry is the product of an optimizer run offline. The second formulation collapses that optimization into the loop.



**Figure 4.8:** Closed loop for learned delay-Doppler shaping. A network reads the range-Doppler map and emits a suppression region  $\Omega_{SL}$  marking the cells to clear, the joint random frequency-modulated waveform optimization of Felton et al. [21] renders the waveform set that pushes range sidelobe modulation out of that region, and the radar transmits it. The mask is the agent’s action, compact and image-shaped, while producing transmittable constant-modulus waveforms is delegated to the optimizer.

Rather than select an entry, the agent emits a few notch parameters (center frequency, depth, and width), and a generator of the same form as the IQWaveNet network from the pre-distortion line of work (section 4.2) renders the corresponding hardware-ready waveform in a single forward pass under 2 ms. The action space stays small while the waveform is continuous rather than the nearest point on a discrete grid. Composing the generator with the pre-distortion stage means the notch survives transmission rather than filling back in under amplifier saturation, the failure mode the spectral term of equation (4.1) was built to catch.

### Shaping the delay-Doppler response

The first two formulations work in frequency and suppress a narrow interferer. The third moves to the delay-Doppler plane and to clutter, which is spread rather than narrowband. Felton et al. [21] jointly optimize a set of random frequency-modulated waveforms across a coherent processing interval so that range sidelobe modulation, the noise-like pedestal that pulse-to-pulse waveform diversity smears across slow-time and that buries small targets beside large clutter, is pushed out of a chosen region of the delay-Doppler response. That region is an input to their optimization. They demonstrate writing arbitrary patterns into the sidelobe floor, evidence of precise control, and more usefully clearing a Doppler band to expose movers that the unoptimized waveform set hides, validated open-air against moving vehicles.

We propose to make that region a learned quantity. An image-like network reads the range-Doppler map and outputs a mask over the cells judged to hold clutter rather than target, and that mask is the suppression region the joint optimization then realizes. Figure 4.8 shows the pipeline. The mask keeps the action compact, an image-to-image map at the resolution of the range-Doppler grid, while the production of constant-modulus, transmittable waveforms is delegated to the waveform optimizer. The optimization is gradient-based, so it can run as an inner loop or be distilled into the policy, presenting the world-model training above with a single differentiable action.

### Risks

The primary risk is over-suppression. A policy that notches too broadly, or a mask that marks too much of the delay-Doppler plane, removes real targets along with the interference. A second risk is that a model trained against the fixed-sequence interferer fails to generalize to real interferers of a different spectral shape, for instance a jammer that does not hop cleanly between discrete bins. A third, specific to the world-model agent, is model exploitation, where the policy learns to beat flaws in the learned model that do not exist on the hardware.

The guard is to keep correcting the model with real transmissions wherever the policy starts to favor regions the model has not seen.

All three formulations still tell the radar what to remove, the interferer, the clutter, the region of the delay-Doppler plane to blank, and reward it for the removal. The final line of work drops even that instruction. Instead of naming what to suppress, it asks the radar to discover from the returns alone which waveform properties make a target easiest to see.

## 4.4. Target-Matched Illumination

**Hypothesis.** A radar can discover from its own returns which waveform properties best reveal a target, without being handed a model of that target in advance. The first experiment narrows this to a single target class, the drone. A learned waveform-selection policy exceeds the standard LFM chirp on drone micro-Doppler signal-to-clutter-plus-noise ratio, the quantity that matters for telling a drone apart from clutter rather than merely detecting it. Pretraining on opportunistic targets (windmills, vehicles, aircraft) is the route there, and the secondary claim is that it reaches the LFM level in fewer coherent processing intervals than a policy starting from scratch. Transfer from opportunistic pretraining to drone returns is not guaranteed, and the experiment is designed to measure whether it happens and at what cost.

### 4.4.1. Problem and Prior Work

Classical theory answers the question this section asks, provided the target is known. Bell [5] notes that “there is no general theorem that maximum output signal-to-noise ratio insures maximum gain of information,” and derives the waveform that does instead: a water-filling solution that maximizes the mutual information between a random target’s impulse response and the received signal, pouring transmit energy into the spectral bands where the target ensemble varies most relative to the noise floor. That result gives target-matched illumination its name, and its message is the premise of this line of work. The best waveform is a function of the target and the clutter, not a universal constant. The LFM chirp dominates fielded systems precisely because it is matched to nothing in particular and well behaved against nearly everything (section 1.2.5).

Applied to surveillance, the theory is circular. Computing the matched waveform requires the target impulse response, or at least its ensemble statistics, yet the target is the thing a surveillance radar does not know. For drones the circularity bites hardest. The discriminating signature is rotor micro-Doppler [3], spread across many Doppler bins, aliased or revealed by the choice of pulse repetition interval, sharpened or buried by bandwidth and polarization, and different for every airframe and aspect angle. The quantity needed to compute the waveform is the quantity the radar is trying to measure.

The established response to this circularity is to close the loop around an estimator. Goodman et al. [64] interleave matched illumination with sequential hypothesis testing, customizing each transmission to the returns from the previous ones, and Bae and Goodman [6] extend the scheme to the constant-modulus waveforms a saturated amplifier demands. Herr et al. [65] carry the information-theoretic objective onto hardware-amenable polyphase-coded FM under the adaptive-on-transmit banner. These works share one structural as-

sumption. Each trusts a parametric model of the target or channel, estimates that model’s sufficient statistics from returns, and computes the next waveform from the estimate. Adaptation ends when the estimate converges, the “learn then deploy” pattern critiqued in section 2.4.

The trusted model is what this setting does not offer. Ground clutter is structured but site-specific, drone signatures vary across airframes and flight states faster than a parametric library can track, and labels accurate enough to supervise the estimate are unavailable (section 3.4). Waveform selection is evaluated almost entirely in simulation against idealized target models [16, 38]. What is missing is a formulation in which the map from return history to next waveform is learned rather than derived from a trusted model, and in which the objective is computable from the radar’s own data stream rather than from a label. Supplying that formulation, and validating it on hardware, is the aim of this line of work.

#### 4.4.2. Framework

We pose illumination as a sequential decision problem over the radar’s action-observation history. At coherent processing interval  $m$  the radar holds the history  $h_m = (a_0, Y_0, \dots, a_{m-1}, Y_{m-1})$  of every waveform transmitted and every range-Doppler map received so far, selects the next waveform  $a_m$  from an action space  $\mathcal{A}$ , transmits it, and computes the range-Doppler map  $Y_m$  defined by equation (1.7) from the return. A policy  $\pi$  maps the history to that choice, and learning maximizes the illumination objective

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{\pi} \left[ \frac{1}{M} \sum_{m=0}^{M-1} r(h_{m+1}) \right], \quad (4.3)$$

where the reward  $r$  is computable from the history alone and  $M$  is the deployment horizon in intervals. The formulation is deliberately a history process rather than a MDP [40, 41]. The scene around a surveillance radar is larger than any state vector the agent could carry, so the Markov assumption criticized in section 2.4 is not available, and the reward must be defined without inspecting a hidden state.

The illumination objective in equation (4.3) contains the prior formulations as special cases, which is what makes it a framework rather than another point design. With a known target ensemble, a single-interval horizon, and mutual information as the reward, the optimal policy ignores the history and its one action is Bell’s water-filling waveform [5]. With a finite hypothesis set and classification confidence as the reward, it reduces to the sequential hypothesis testing scheme of Goodman et al. [64]. Sense and Notch (section 4.3) is the same loop with a reward that names what to remove from the return. Target-matched illumination removes that last instruction. Nothing names the target, so the reward must be constructed from observables and the policy must be learned from interaction.

The choice of reward is therefore the central open question of this line of work, not an implementation detail. A reward that asks whether a waveform revealed the target cannot be a function of a single return, because target presence is the question. It must be a functional of the history, scoring the structure that real movers create across consecutive returns and that clutter and noise do not. Several observable rewards are worth trying, each an estimator of the same underlying quantity, how much the return history reveals about the movers in the scene.

**Track persistence** Reward a range-Doppler cell whose energy stays coherent along a physically plausible track over consecutive intervals. A real mover produces a temporally consistent signature, clutter and thermal noise do not, and scoring consistency needs no label.

**Surprise against a scene model** Reward the waveform whose return departs most from the world model’s prediction of the static background, concentrated in a small moving region. The world model already supplies that prediction.

**Micro-Doppler structure** Reward returns whose Doppler spectrum shows the harmonically spaced sidebands of rotor modulation. This proxy is double-edged, since a wind turbine produces similar structure that the experiment would have to exclude.

**Downstream information gain** Reward the rise in a self-supervised classifier’s confidence, or the mutual information between the waveform choice and the resulting return, neither of which needs ground truth.

None of these is obviously correct, and **finding a reward that survives contact with real clutter is a research question** in its own right. While that question stays open, a genie reward scaffolds the experiments. Ground truth from a GPS tracker labels the target’s range bin, and the measured signal-to-clutter-plus-noise ratio in that bin is the reward. This is genie-assisted in the same sense as the genie notcher of section 4.3, a fine training signal for a controlled experiment that presumes knowledge of the very thing an operational radar is trying to find.

Whichever reward trains the policy, performance is reported against one fixed, externally meaningful metric. The primary metric throughout is the signal-to-clutter-plus-noise ratio integrated over the micro-Doppler band in the target range cell  $n_T$ ,

$$\text{SCNR} = \frac{\sum_{d \in \mathcal{B}} |Y[n_T, d]|^2}{|\mathcal{B}| \hat{\sigma}_{n_T}^2}, \quad (4.4)$$

where  $Y[n, d]$  is the range-Doppler map defined in equation (1.7),  $\mathcal{B}$  is the set of Doppler bins spanning the micro-Doppler sidebands, and  $\hat{\sigma}_{n_T}^2$  is estimated from a ring of guard and reference cells around the target, following the standard cell-averaging constant-false-alarm-rate scheme [7]. This estimate assumes range-stationary clutter, a standard simplification that is expected to hold in the target range cell for the low-PRF operation of PARSAX. The set  $\mathcal{B}$  is defined a priori as the Doppler bins within  $\pm 200$  Hz of the body Doppler estimated from the GPS track, covering the expected rotor sideband extent. A single delay-Doppler cell cannot hold a rotor signature, which spreads across many Doppler bins, so the metric integrates over the band rather than a point.

The action space  $\mathcal{A}$  grows in 3 rungs, each validated before the next. The first rung is a binary choice, the standard LFM chirp against one alternative of different pulse repetition interval or bandwidth, and exists to establish that waveform choice produces a measurable difference at all. The second rung is a library pre-generated across pulse repetition intervals (3–5 values), bandwidths (3–5 values), and polarization configurations (2 values), an order 30–50 discrete set. The final rung relaxes the grid. Rather than index a fixed library, the policy emits waveform parameters that an IQWaveNet-style generator (section 4.2) renders into a transmittable waveform, so the agent designs continuously along the same axes

instead of snapping to the nearest precomputed point. The agent's input at each coherent processing interval is a compact descriptor of the returned range-Doppler map. It combines micro-Doppler energy across the Doppler bins in the target range cell (tentatively  $\sim 64$  bins) with a clutter figure of merit derived from the surrounding cells (a  $\sim 16$ -element context vector).

The world-model agent of section 4.3 (figure 4.7) is the proposed vehicle for optimizing the illumination objective. The encoder, latent dynamics, and policy carry over unchanged, and only the reward changes, from suppressing a named interferer to revealing an unnamed target. A model of the scene and hardware, learned from logged returns, lets the policy train in imagination and spend real transmissions only where the model is wrong.

### 4.4.3. Simulation

Simulation plays a different role here than in the pre-distortion line of work, where it stood in for the amplifier until hardware was available. Here, simulation is the one setting where the matched waveform is computable, because the experimenter knows the target impulse response that the radar does not. A simulated scene of point scatterers and rotating rotor blades, following the micro-Doppler model of Chen et al. [3], placed in structured clutter, therefore supports 3 screens that hardware cannot. The first screen asks whether per-interval SCNR differences across the action space exceed measurement noise at all. If no waveform in the library separates from the rest, no policy can learn anything, and the question is settled before hardware time is spent. The second asks whether a policy trained with the genie reward recovers the computable optimum, which validates the learning machinery against theory rather than against another learned system. The third asks which of the observable rewards track the genie reward across scenes. Only the proxies that survive this screen graduate to hardware.

The world model then bridges simulation and hardware. Trained first on simulated returns and then on logged opportunistic returns from PARSAX, it lets the policy continue training in imagination on data with real-hardware statistics before any drone is in the air, the same pattern section 4.3.3 uses for the interference problem.

### 4.4.4. Hardware Implementation

The framework asks little of the hardware. It needs a transmitter that accepts an arbitrary waveform every coherent processing interval, a receiver that returns the corresponding range-Doppler map, and a round trip faster than the scene changes, exactly the loop section 4.1 closes. Any sufficiently capable software-defined radar satisfies this. PARSAX is the instance available to this project, and 2 of its properties make it an unusually good one. It transmits and receives on orthogonal polarizations at the same time, so polarization enters the action space directly, a discriminant the micro-Doppler-only literature leaves on the table, and one along which a drone body and its rotors are expected to differ from ground clutter. Its surroundings, windmills, passing vehicles, and general aviation approaching Rotterdam-The-Hague Airport, supply opportunistic targets with distinct micro-Doppler statistics and clutter contexts, a free source of training signal before any drone is in the air. Pretraining on targets of opportunity is not a new idea in principle, but to our knowledge it has not been demonstrated in hardware for the radar waveform selection problem.

The first hardware milestone is the binary rung of the action space. 2 waveforms switch every coherent processing interval against reference targets, and a paired z-test on the per-interval SCNR difference checks whether waveform choice is detectable above measurement noise on real returns. If the difference between the 2 waveforms is not detectable on real targets, the full library design has no foundation, and the simulation screens must be revisited before anything larger is attempted.

Training begins on targets of opportunity before any drone is scheduled, and the policy updates continuously through this phase, converging toward waveform preferences that improve micro-Doppler signal-to-clutter-plus-noise ratio across a range of target types. When the drone campaign begins, the policy carries its pretrained weights and continues updating online. At each coherent processing interval the policy selects or generates a waveform, PARSAX transmits and receives it, the range-Doppler map is processed, and the GPS tracker labels the drone's range bin (section 3.4) for the genie reward and for evaluation. The reward is computed and accumulated, and PPO updates run periodically through the flight. The policy does not reset between sessions. Experience from one flight informs the next, and flying multiple drone types across the campaign acts as a natural curriculum.

Three baselines bracket the result. Random waveform selection sets a floor, the standard PARSAX LFM waveform represents current practice, and a post-hoc oracle picks the best library waveform after the fact given the full range-Doppler map and the GPS-confirmed drone location. The oracle bounds the library, not what the generative rung can reach, so the final rung is scored against the LFM baseline and the library oracle together. Two quantities are reported together.

**Convergence rate** How many coherent processing intervals the online policy needs to reach within a fixed margin of the oracle.

**Steady-state SCNR** The performance the policy holds once converged, relative to the oracle and the LFM baseline.

Convergence speed without steady-state quality is not a useful result.

### Risks

Reward collapse is the most likely failure. The pretraining-to-drone transfer is not guaranteed. Windmill and vehicle pretraining may reinforce waveforms that degrade drone micro-Doppler, for instance a pulse repetition interval that aliases the rotor rotation frequency or a bandwidth that places windmill clutter in the same range-Doppler bins as the drone. A further risk is reward hacking, where the policy maximizes an observable proxy by chasing clutter that mimics target structure, with the wind turbine as the standing example. The simulation screen reduces this risk but cannot eliminate it, since the proxies that fail will fail against precisely the clutter the simulation did not contain.

# 5

## Preliminary PARSAX Experiments

Two measurement campaigns establish that PARSAX can be operated for adaptive waveform work and expose the one engineering constraint that blocks rapid iteration: the FPGA firmware must be rebuilt for each new pulse length. PARSAX is well documented, but the gap between documentation and being able to operate it for adaptive waveform work is substantial at this stage of the PhD. Even more so trying to operate it in quite a different manner than we are accustomed has surfaced many engineering challenges. Uploading arbitrary waveforms programmatically, triggering captures reliably, and moving data off the hardware fast enough for closed-loop processing all require working through issues the current documentation does not cover.

This chapter records the first year of closing that gap. It covers what the hardware actually looks like from a software perspective, what two measurement campaigns revealed about its capabilities and constraints, and what still needs to change before any of the experiments in chapter 4 can run.

Notably, the three experiments described in chapter 4 share a hardware precondition that no simulation can satisfy. **Learned Digital Pre-Distortion** trains a network to invert a physical power amplifier: the thermal drift and saturation behavior that motivate the approach only appear in a real device running at power. **Sense and Notch** must close its decision loop within the pulse repetition interval, and whether the inference-and-upload cycle fits that budget is a question only real hardware can answer. **Target-Matched Illumination** pretrains on opportunistic returns from vehicles, windmills, and aircraft before any drone is scheduled, and simulated clutter statistics are proxies for real ones, not the thing itself.

This chapter presents the 2 measurement campaigns that establish whether PARSAX can be operated in the way those experiments require. The first campaign, in October 2025, flies a drone at 500–600 m and uploads 25 waveforms by hand through the AWG front panel. The second, in May 2026 at Poeldijk, flies at 9 km with waveforms uploaded via a scripted interface. Together they demonstrate that automated waveform upload, dual-polarimetric capture, and on-site processing all work. They also expose the one engineering constraint that blocks rapid iteration. The FPGA firmware ties pulse length to the processing pipeline,

requiring a bitstream rebuild for each new waveform duration. That firmware modification is the initial deliverable described in section 4.1.

## 5.1. Hardware Chain

PARSAX transmits and receives on 2 orthogonal polarizations simultaneously [66].<sup>1</sup> On transmit, a Tektronix AWG5014C AWG [67] produces the baseband signal as 14-bit samples at up to 1200 MP/s, centered at an intermediate frequency of 125 MHz. An upconverter shifts this to a carrier at 3315 MHz before the power amplifier and transmit antenna. The standard waveform is an LFM chirp sweeping from 100 MHz to 150 MHz at intermediate frequency (3290 MHz to 3340 MHz over the air), 0.050 ms duration, repeated with a pulse repetition interval of 0.200 ms. Bandwidth is nominally 50 MHz, with some headroom to 100 MHz before hardware non-linearities become significant.

On receive, each polarization channel goes through an analog mixer, then into a 2-channel Texas Instruments ADS54J60 ADC producing 16-bit real samples at 500 MP/s. These feed a Xilinx Virtex-7 690T FPGA on a Vadatech PCI516 card, which streams data to the host PC over PCIe Gen3  $\times 8$  (of which, Gen3  $\times 4$  is available for direct memory access). The FPGA currently runs a 16 384-point fast Fourier transform and multi-stage decimation before the data ever reaches the host, which is where things get complicated (see section 5.2).

One of the first things needed to reason about GPU-first processing is a clear picture of what connects to what and how much data each interface can carry. Figure 5.1 is the result of that exercise. It shows every component in the signal chain, the bus between each pair, and the bandwidth available at each stage. The exact numbers matter because whether all matched filtering and down-conversion can move from the FPGA to the GPU is ultimately a bandwidth question. This remains the biggest outstanding piece of work to solve before moving this work forward.

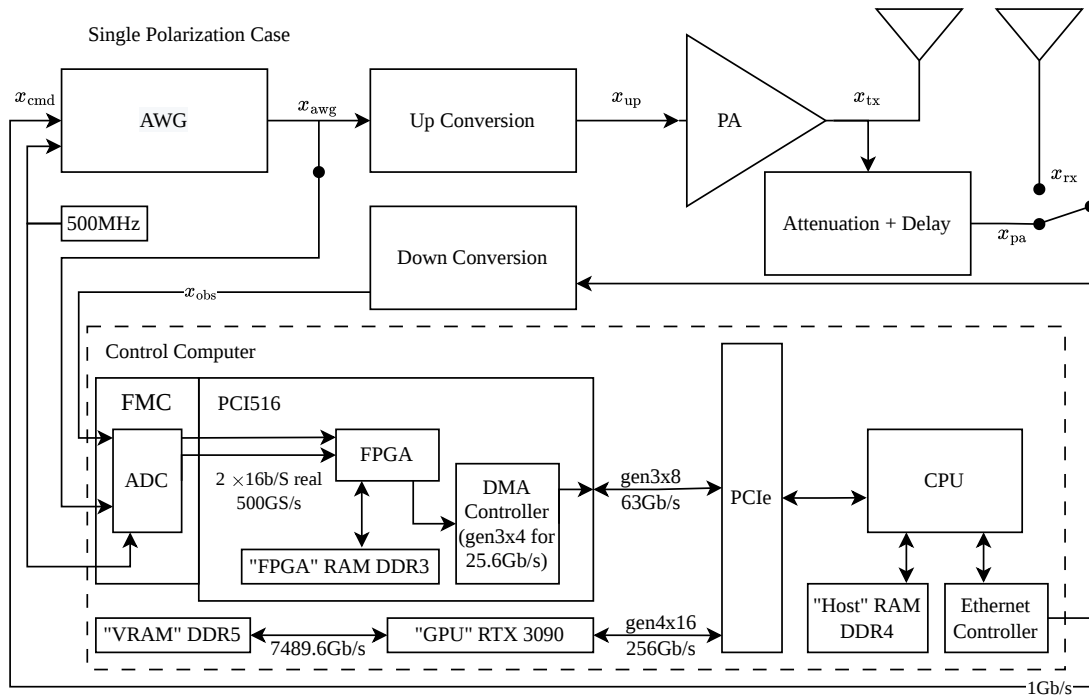
## 5.2. Preliminary Experiments

The first practical issue is waveform upload. The AWG is nominally controllable over TCP (Transmission Control Protocol), but in practice not yet from the Ubuntu machine. Waveforms are currently stored as normalized floating-point `.txt` files, with newline-separated strings, 6–8 decimal places, and one sample per line. A single 0.050 ms waveform at 400 MP/s is 20 000 samples, in this format that is about 39 MB of text for what is really 40 kB of data. The AWG’s own manual describes a native 14-bit integer format of 2 bytes per sample, loaded directly into hardware memory. At 1000 Mb/s Ethernet, a 40 kB waveform transfers in under a millisecond, the 39 MB text version takes about a third of a second, before conversion overhead on the hardware side.

In October 2025, the group started a project with the Dutch Ministry of Defence on drone detection using PARSAX. The goal is to run initial drone experiments to establish a baseline before any adaptive waveform work. The primary target is a DJI Matrice 600 drone flying at X-Fields on campus, about 500–550 m from the radar at about 56 m above ground.

---

<sup>1</sup>Polarization is the orientation of the electric field of a transmitted wave. Targets scatter differently depending on the alignment between the wave’s electric field and the target’s geometry. A dual-polarimetric radar transmits and receives on 2 orthogonal polarization channels simultaneously.



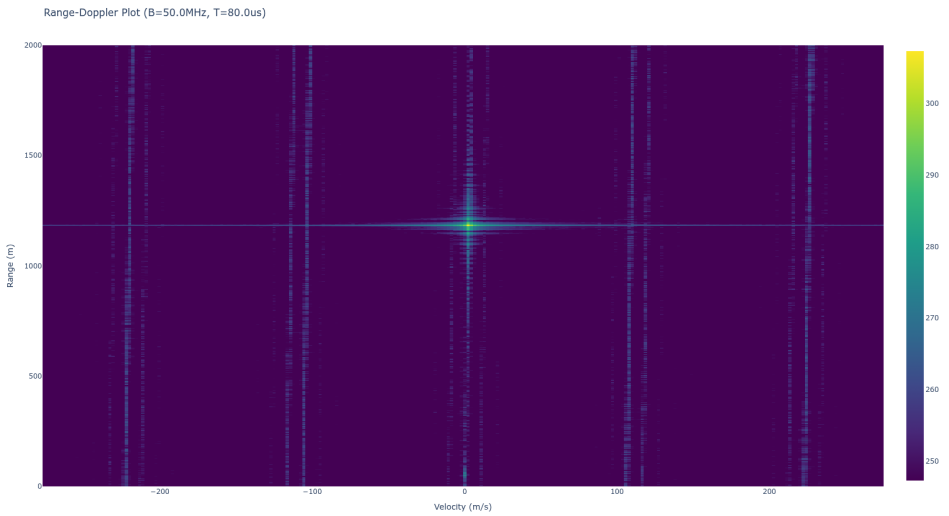
**Figure 5.1:** Signal flow through the PARSAX hardware chain, single-polarization, with the bandwidth-relevant interfaces labeled. A 14-bit integer sample array travels from the host over Ethernet to the AWG, through the power amplifier and a directional coupler that taps a loopback monitor signal, and onward to the transmit antenna. The receive path digitizes at 500 MP/s per polarization channel and returns samples to the host via PCIe, where they are available to the GPU on the same bus.

The five waveform durations are a geometric progression by powers of two in sample count at 1200 MP/s: 0.013 ms ( $2^{14}$  samples), 0.027 ms ( $2^{15}$ ), 0.050 ms ( $2^{16}$ , matching the reference waveform in earlier work [17]), 0.100 ms ( $2^{17}$ ), and 0.200 ms ( $2^{18}$ ). The five bandwidths are 12.5, 25, 50, 75, and 100 MHz, ranging deliberately above the 50 MHz design specification to see where things break.


Working through the practical constraints reveals several things. The AWG's Repeat and Goto commands handle pulse repetition internally, so repeated waveform copies are unnecessary. Different-length sequences on the horizontal and vertical channels simultaneously do not work. The old "snap" acquisition system reliably records about 500 ms, enough to capture the full 25-waveform sweep. We upload the 25 waveforms manually through the AWG front panel. That is already tedious, it also contains at least one incorrect jump sequence that corrupts part of the data. At any larger number it is untenable.



**Figure 5.2:** Aerial view of the experimental area. PARSAX radar at the origin, chimney target at 1180 m, X-Fields at 600 m.



**Figure 5.3:** Range-Doppler map from the first 2025 PARSAX campaign (50 MHz bandwidth, 0.080 ms sweep time, 25-waveform sweep uploaded manually). The chimney return at 1180 m with near-zero Doppler confirms that the custom waveform was transmitted and returned coherently. This baseline establishes that waveform upload and range-Doppler processing work before any automation is introduced.



**Figure 5.4:** Range-Doppler plot showing a timing anomaly. The ground return and static chimney wander across range bins. The source remains unresolved, candidates include the long-pulse capture settings, the maximum sample rate, the multi-waveform packing scheme, and a newer capture firmware reproduces the same anomaly with the standard 1 ms sweep at a more modest sample rate.

Sethuraman et al. [68] previously use PARSAX for drone payload classification, flying a 4-Rotor and 6-Rotor drone at 570 m on rectangular and linear paths with a 50 MHz LFM waveform, achieving 90–95% classification accuracy. That experiment extracts polarimetric features from the scattering matrix, ours does not attempt classification. We are learning the system and establishing which waveform parameters are feasible before introducing any adaptive waveform logic.

The ability to upload custom waveforms to the AWG is well documented, and the existing system uses this interface, but no reliable automated demonstration exists. Uploading just 25 waveforms during the first campaign was highly manual and error-prone.

The FPGA currently runs a full fast Fourier transform and decimation pipeline, which means the data format and timing assumptions baked into downstream software are tied to specific pulse lengths. Changing the pulse duration is not a matter of uploading a new waveform, it requires rebuilding the FPGA bitstream. This blocks rapid iteration. We are working on a firmware modification to allow waveform agnostic processing, by moving nearly everything other than the initial downmixing to GPU. If the firmware path stalls, the fallback is to lock the pulse length at  $2^{16}$  samples and offload processing to the GPU within that fixed frame, less flexible, but sufficient to close the loop and begin Learned Digital Pre-Distortion. An earlier exploration of Hardcaml<sup>2</sup> for rapid FPGA iteration was explored but abandoned. Build times alone prevent anything resembling real-time iteration.

Two channels at 500 MP/s is a high data rate, but an order-of-magnitude less than the memory bandwidth of the Nvidia RTX 3080. The PCIe interface is not the limiting factor. Using the FPGA for technique-agnostic work (maximizing PCIe throughput, arranging raw samples into maximum packet sizes, enabling direct memory access by the GPU) while moving all signal processing to the GPU would decouple the processing from the waveform

<sup>2</sup><https://github.com/janestreet/hardcaml/>

parameters. Rebuilding a GPU kernel for a new pulse length takes minutes. Rebuilding an FPGA bitstream takes hours and requires proprietary tooling. These are the two issues described earlier in section 4.1 is meant to fix.

## 5.3. Second Campaign: Poeldijk, 9 km Baseline (2026)

The first experiment flies at 500–600 m with waveforms uploaded by hand through the AWG front panel. The second campaign targets a drone at roughly 9 km over open terrain at Poeldijk, matching the geometry of a 2018 PARSAX measurement that produces strong micro-Doppler signatures on a DJI Matrice 600 6-Rotor drone. That archived dataset is the reference, a comparable result from the rebuilt system confirms the hardware chain works at this range before any adaptive waveform logic is introduced.

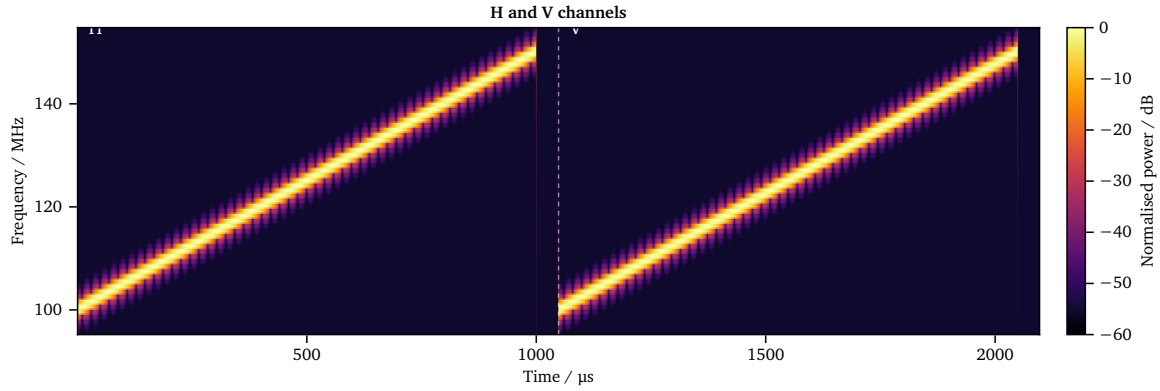
### 5.3.1. Waveform Design

The campaign uses 2 prepared waveform configurations. Iterating with Prof. Oleg Krasnov, we selected the following 2 waveforms.


**Option 1: time-multiplexed polarimetric.** Horizontal and vertical polarizations transmit in sequence within a single  $2^{20}$ -sample pulse repetition interval of 2097  $\mu\text{s}$ . Horizontal sweeps 0–1000  $\mu\text{s}$ , vertical sweeps 1049–2049  $\mu\text{s}$ , each with roughly 49  $\mu\text{s}$  of silence before the other channel begins. Cross-polarization leakage is eliminated by construction, the channels do not overlap. At 50 MHz bandwidth this matches the 2018 sweep duration exactly, enabling a direct comparison.

**Option 2: staggered simultaneous.** Both polarizations sweep the same 50–100 MHz band simultaneously, but the vertical channel starts 80  $\mu\text{s}$  after horizontal. Horizontal sweeps 1–181  $\mu\text{s}$ , vertical sweeps 81–261  $\mu\text{s}$ , both share a pulse repetition interval of  $2^{17}$  samples (262  $\mu\text{s}$  at 500 MP/s), roughly 8 times shorter than Option 1.

The 80  $\mu\text{s}$  stagger pushes cross-polarization leakage beyond the target range. A stagger exactly equal to the 60  $\mu\text{s}$  round-trip time to 9 km would place vertical transmit leakage precisely on the 9 km range bin of the horizontal receive channel. Setting the stagger to 80  $\mu\text{s}$  moves the leakage to  $c \times 80 \mu\text{s}/2 = 12 \text{ km}$ , 3 km beyond the drone hover point, where it is harmless.



**Figure 5.5:** Spectrogram of Option 1 at 50 MHz bandwidth. Horizontal polarization (H) sweeps 100–150 MHz over 0–1000  $\mu\text{s}$ , vertical (V) begins at 1049  $\mu\text{s}$  and sweeps the same band to 2049  $\mu\text{s}$ . The 2 chirps are non-overlapping by construction, eliminating cross-polarization leakage. The full pulse repetition interval is  $2^{20}$  samples (2097  $\mu\text{s}$  at 500 MP/s).



**Figure 5.6:** Spectrogram of Option 2 at 50 MHz bandwidth, horizontal (top) and vertical (bottom) channels shown separately. Both polarizations sweep 100–150 MHz, the vertical channel starts 80  $\mu\text{s}$  after horizontal, visible as the delayed chirp onset in the lower panel. The stagger places cross-polarization leakage at 12 km, 3 km beyond the 9 km drone hover point. The full pulse repetition interval is  $2^{17}$  samples (262  $\mu\text{s}$  at 500 MP/s), roughly 8 times shorter than Option 1.

The short pulse repetition interval of Option 2 has practical consequences. With  $2^{14}$  pulses, the coherent processing interval is 4.3 s, against 34.4 s for Option 1 at the same pulse count. Running 4 bandwidths (25, 50, 75, and 100 MHz) with Option 2 takes under 20 s total, the equivalent sweep with time-multiplexed transmission takes over 2 minutes. This speed allows a full bandwidth comparison within a single hover, before the drone drifts or the battery changes.

### 5.3.2. Tooling

All waveforms are precomputed and stored as 14-bit `.bin` files before the campaign day. Technicians have developed libraries that let us interact with PARSAX much as we would with a software-defined surrogate. Upload and sequencing run via `awg5014_python`, a single Python script uploads all configurations to the AWG and programs the GOTO transitions between them. Capture runs via `ptshm-python`. The experiment directory is self-contained and starts with `uv sync && uv run run_experiment.py`.

This is the first field use of `awg5014_python`. The first campaign requires an afternoon of front-panel interaction to upload 25 waveforms, and still contains at least 1 incorrect jump sequence that corrupts part of the data. This campaign programs 5 configurations in seconds, the script is the verbatim record of what the AWG executes. The same infrastructure should scale, with little or no modification, to the hundreds of waveform configurations the proposed adaptive waveform experiments require.

Running this campaign before the adaptive loop is online is deliberate. Designing the staggered simultaneous waveform, coordinating the airspace logistics, building and testing the scripted upload, and executing the campaign with the measurement group are not incidental preparation, they are evidence that the platform can be operated in the way the proposed experiments require. The difference between a simulation result and a proposal built on real measurements is a flight campaign.

### 5.3.3. Results

On 21 May 2026 we collected 96 GB of dual-polarimetric returns: a static chimney at 1.17 km for calibration and the A13 motorway at 3–10 km azimuth  $150^\circ$  as a moving-target scene. Both scenes use Option 2 at 50 MHz bandwidth.

Processing follows a stretch-processing receiver. The transmit monitor (HT) and receive (HR) channels are loaded as raw complex ADC samples and the beat signal is formed as

$$s[n, m] = H_R[n, m] \cdot H_T^*[n, m], \quad (5.1)$$

where  $n$  indexes slow time and  $m$  indexes fast time. Conjugating the transmit reference sets the sign convention for a down-chirp: without it, targets land at negative range bins.


The Option 2 chirp occupies  $N_{\text{chirp}} = 11\,250$  of the 16 384 ADC samples per sweep ( $180\ \mu\text{s}$  of the  $262\ \mu\text{s}$  pulse repetition interval). Using  $c/2B$  as the range bin spacing gives 3 m per bin. The correct value accounts for the duty factor,

$$\Delta R = \frac{c N_{\text{chirp}}}{2 B N_{\text{bins}}} \approx 2.06\ \text{m}, \quad (5.2)$$

and the range window is applied only over the  $N_{\text{chirp}}$  active samples. A 512-sweep coherent processing interval ( $\approx 134\ \text{ms}$ ) slides across the 16 384 sweeps. Slow-time mean subtraction per range bin removes the dominant static clutter before the Doppler FFT, and power is expressed as SNR relative to the per-range-bin noise floor estimated from Doppler bins with  $|v| > 5\ \text{m/s}$ .

Figure 5.7 shows a representative HH range-Doppler map from the highway scene. Vehicle returns are visible from 3 to 9 km at velocities consistent with motorway traffic. Both

approach and recede directions are populated. The residual feature near  $\pm 40$  m/s is a Doppler sidelobe of a bright static reflector at  $\approx 13.5$  km, reduced but not eliminated by the Blackman range window.



**Figure 5.7:** HH range-Doppler map of the A13 motorway, 21 May 2026 (50 MHz bandwidth, Option 2, 512-sweep coherent processing interval,  $\approx 134$  ms). *Left:* aerial view aligned to the range axis. *Right:* power relative to the per-bin noise floor; blue lines bound the named scene regions shown at right. Positive velocity is radially away from the radar. Static clutter at 0 m/s is suppressed by slow-time mean subtraction. Some aliasing artifacts have been traced back to now rectified issues with filtering before decimation in receive pipeline.

Running the Python script to configure and upload the waveforms takes 1.7 s, faster than the 4.3 s capture. The 1.7 s upload latency is roughly 6500 times the 0.262 ms pulse repetition interval, ruling out pulse-to-pulse adaptation, but well within the budget for coherent-processing-interval-to-coherent-processing-interval adaptation at multi-second cadence.

## 5.4. Relation to the Proposed Lines of Work

PARSAX hardware data are essential to each of the 3 proposed lines of work described in chapter 4, for different reasons.

Learned Digital Pre-Distortion, and Sense and Notch, specifically addresses the gap between a simulation and the physical system. The pre-distortion experiment's central claim is that a network trained on loopback measurements tracks the real amplifier as it heats and changes state. In the sense and notch experiment, the claim hinges on being able to consume and learn from real returns. A simulation is insufficient to prove these claims.

Target-Matched Illumination pretrains on opportunistic targets: windmills, vehicles, and aircraft. Those targets are in the PARSAX field of view continuously, and the pretraining signal is free, but only if the radar is collecting real returns. The drone experiment then claims that pretraining on opportunistic targets transfers to drone micro-Doppler. That transfer must be tested against real drone returns at range. Simulated micro-Doppler from a rotating rotor at 9 km, including ground clutter statistics and multipath at that geometry, is not a reliable substitute for the real thing at the level of fidelity that determines whether transfer occurs or not.

There is a fourth reason that spans all 3 experiments and connects to the continual learning gap (section 3.3). A hardware-in-the-loop radar generates a data stream whose statistics depend on the waveform being transmitted. The radar is a participant in its own environment. This closed-loop dependency is what makes continual learning on radar hard, and it is the property that offline replay datasets cannot replicate. In replay, the environment is fixed and the agent is a passive observer of recorded history. On hardware, transmitting a different waveform changes what comes back. They are evidence that the platform generates the kind of data those experiments require, at the rates and volumes they require it, and that the infrastructure to do so reliably is in place.

# 6

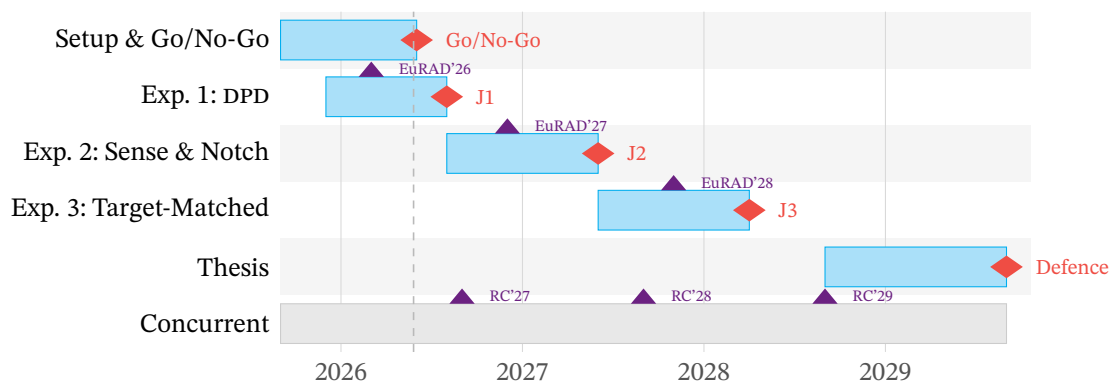
## Administrative Details

A research plan, doctoral education plan, and data management plan, targeting a PhD defense no later than September 2029.

### 6.1. Research Plan

Phase	Period	Key Deliverable
Setup and Go/No-Go	Sep 2025 – Jun 2026	Go/No-Go defense
Experiment 1: Digital Pre-Distortion	Dec 2025 – Aug 2026	Journal 1 submission
Experiment 2: Sense and Notch	Aug 2026 – Jun 2027	Journal 2 submission
Experiment 3: Matched Illumination	Jun 2027 – Apr 2028	Journal 3 submission
Thesis	Sep 2028 – Sep 2029	Defense

**Table 6.1:** PhD research arc by phase (timelines may shift as the work develops).



**Figure 6.1:** Planned research timeline (May 2026). DPD runs 8 months overlapping the Setup phase, Sense and Notch and Target-Matched Illumination run 10 months each, thesis 12 months. The gap between J3 and thesis start is intentional: journal revisions and resubmissions land there and fold into the Concurrent row. Blue bars show phases, red diamonds mark key deliverables (Go/No-Go, journal submissions J1–J3, and defense), purple triangles mark conference submission dates: one per experiment row, plus an IEEE Radar Conference submission each fall in the Concurrent row (conferences themselves the following spring). Dashed line is today.

The PhD targets one journal paper per experiment, with each submission marking the close of its phase. Given the applied nature of this work, findings aim to be published in IEEE Transactions on Aerospace and Electronic Systems, IEEE Transactions on Radar Systems, or IET Radar, Sonar, and Navigation. If some concurrent projects result in testbeds or open sourcing software, these may be submitted as tutorials in AESS Magazine. EuRAD 2026 (digital pre-distortion) has already been accepted [58].

## 6.2. Doctoral Education Plan

Delft University of Technology requires 45 credits organized into three categories: Transferable, Discipline, and Research skills, 15 credits each. Transferable skills have been intentionally front-loaded, discipline skills slightly so, with the intention of increasing research velocity once I am more comfortable with the problem space.

### 6.2.1. Transferable (15/15)

Code	Course Title	Credits	Status
T4.G1-A	PhD Start-up Module A	1.5	Completed
T4.G1-B	PhD Start-up Module B: Scientific Integrity	0.5	Completed
T1.C1	Scientific Storytelling	2	Completed
T1.A7	Data Visualization	1	Completed
T1.A1	Presenting with Confidence	3	Completed
T1.A6	Voice Training	1.5	Completed
T2.C1	Effective Negotiation	2.5	Completed
T4.G5	Career Development: Personal Branding	1	Completed
T4.B5	Project Management for PhD Candidates	2	Completed

**Table 6.2:** Transferable credits (15 completed).

All completion certificates are uploaded to <https://dma.tudelft.nl>.

### 6.2.2. Discipline (0/15)

The Luxembourg summer school on Integrated Sensing and Communication (ISAC) was selected for its speaker overlap with the literature this PhD addresses. Most of the group attends the International Summer School on Radar/Synthetic Aperture Radar (SAR), the Bonn school may shift summer 2027, or be replaced with other relevant masters-level classes at Delft University of Technology. The remaining 6 credits will be filled from courses not yet finalized, 2 strong candidates are a graduate-level machine learning course through the 4TU network (2–3 credits) and participation with an oral presentation at IEEE Radar Conference (2–3 credits).

Course	Credits	Status
2026 Summer School on ISAC, Luxembourg	4	Planned
18th International Summer School on Radar/SAR, Bonn	5	Planned
Master level course (Perhaps Radar I)	5	Planned
Short course at some radar conference	1	Planned

**Table 6.3:** Discipline-specific credits (target 15).

### 6.2.3. Research (>4/15 Completed)

Research credits accrue through output and service alongside the main experiment arc. Completed so far:

- Submitted a first-author conference paper to EuRAD 2026 [58] (2–3 credits, depending on acceptance outcome).
- Participated as instructor in 2 Software Carpentry workshops at Delft University of Technology (2 credits).
- Acted as reviewer for at least 1 IEEE conference submission (credit pending logging).

Planned additions to reach 15 credits:

- “Participation in work consultation with research partners” – Spend some time in person at Thales Germany in Ditzingen.
- “(International) internship of at least 1 month with another institute” – Visit University College London (Prof. M. Ritchie) or University of Kansas (Prof. S. Blunt) to work with different hardware.
- “Teaching Assistance: Preparing and giving a lecture in regular MSc course”
- “Supervising a (MSc and/or BSc) student”

## 6.3. Concurrent Activities

Viktor Vozár (a fellow PhD candidate implementing PPO waveform selection in simulation) is running a separate PARSAX experiment in collaboration with this project, giving low-stakes early platform exposure under someone else’s experimental design. The Netherlands Ministry of Defence (MoD) dual-polarimetric drone data collection, begun in autumn 2025, is coordinated from within this project and produces a dataset the MS<sup>3</sup> group can use independently.

## 6.4. Data Management Plan

No industrial proprietary data has been used at this stage. All data is generated by our own simulators or sensors. A discussion with the faculty data steward will be planned in the future as needed.

# 7

## Self-Reflection

Delft is so cute, and the group at MS<sup>3</sup> has been welcoming. Prof Fioranelli and Prof Yarovy have both been generous with their time and supportive of an approach that has involved a fair amount of wandering and tinkering with hardware.

The sponsor, Thales Germany, has been generous with scope. The brief was broad: waveform diversity for ground surveillance radar. In some ways that is ideal, with nobody constraining curiosity, but it also leaves me not totally certain what the important problems in this space are.

The roughly unlimited scope of *waveform diversity* is part of the problem. I am trying to let PARSAX somewhat constrain things. This means multiple sensors, an array, moving the sensor, are all reasonable things being ignored. Less reasonable are various history process learning, drone-facing-down, single-use-fiber bistatic, a multi-static network in the Cyber Zoo, and tokenization ideas. These occupy a large share of my imagination.

Hardware has been difficult in a different way than I had anticipated. My mental model was more of an async ROS node or Gym environment. The reality is a sequence of snapshots with overhead and synchronization constraints, and often thinking far more processed. I still do not think in beat frequency or filters at all. As a result I am back to iterating on simulations that leave me frustrated that I can't even create a problem, let alone solve it. All typical I believe though, for a PhD.

Longer term, I try to keep my nose to the ground for macro trends in academic careers, and in fairly technical startups, especially the somewhat recent “dual use” gold rush. This has me wanting to explore things like how large language models may enable different types of research<sup>1</sup>, and also has me wondering if there could be ways to shape my work into something that would make me attractive for counter Drone grants<sup>2</sup>.

Overall, I look forward to moving ahead with this research program.

---

<sup>1</sup><https://github.com/karpathy/autoresearch>

<sup>2</sup><https://www.canada.ca/en/department-national-defence/programs/defence-ideas>

# 8

## Conclusion

Radar waveform design rests on decades of theory, yet fielded ground surveillance radars still do not search that design space using their own returns. The prevailing approach fixes a waveform strategy. Such a system knows how to transmit but has no mechanism for learning what to transmit. This PhD asks whether that gap can be narrowed on real hardware.

The 3 proposed experiments are arranged as a single progression rather than 3 separate results. Learned pre-distortion asks whether a model can learn the hardware's forward map and invert it, with a fully specified target waveform. Sense and Notch removes the specified target and asks the system to respond to an external stimulus in real time. Target-Matched Illumination removes the specification entirely, asking which waveform properties best reveal specific targets from the returns alone. Each step relaxes one assumption, and each depends on the one before it.

A model that survives contact with real hardware carries something the simulation literature, on its own, cannot confer. Real hardware generates data whose statistics reflect genuine amplifier nonlinearities, real clutter, and real targets, properties no simulation reliably replicates at the fidelity that determines whether transfer occurs. The Delft University of Technology radar PARSAX serves as an excellent test platform. The methods depend only on a closed loop between a generated waveform and a captured return, which any sufficiently capable software-defined radar can in principle provide. If the loop closes, the question facing this PhD shifts from whether a waveform-learning loop can be built to what is worth learning once it runs. That, more than any single experiment, is what the remaining years of this work are meant to explore.

# References

- [1] S. D. Blunt and E. L. Mokole, "Overview of radar waveform diversity," *IEEE Aerospace and Electronic Systems Magazine*, vol. 31, no. 11, pp. 2–42, Nov. 2016, ISSN: 0885-8985. DOI: 10.1109/MAES.2016.160071 Accessed: Sep. 18, 2025. [Online]. Available: <http://ieeexplore.ieee.org/document/7771665/>
- [2] M. E. Davis, "A history of battlefield surveillance radar," in *2015 IEEE Radar Conference (RadarCon)*, May 2015, pp. 1345–1350. DOI: 10.1109/RADAR.2015.7131204 Accessed: Oct. 13, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/7131204/>
- [3] V. C. Chen, F. Li, S.-S. Ho, and H. Wechsler, "Micro-Doppler effect in radar: Phenomenon, model, and simulation study," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 42, no. 1, pp. 2–21, Feb. 2006, ISSN: 0018-9251. DOI: 10.1109/TAES.2006.1603402
- [4] *686-2024 - IEEE Standard for Radar Definitions*. IEEE, 2024, ISBN: 979-8-8557-1452-4.
- [5] M. R. Bell, "Information theory and radar waveform design," *IEEE Transactions on Information Theory*, vol. 39, no. 5, pp. 1578–1597, 1993. DOI: 10.1109/18.259642
- [6] J. Bae and N. A. Goodman, "Evaluation of modulus-constrained matched illumination waveforms for target identification," in *2010 IEEE Radar Conference*, May 2010, pp. 871–876. DOI: 10.1109/RADAR.2010.5494497 Accessed: Oct. 13, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/5494497/>
- [7] N. Levanon and E. Mozeson, *Radar Signals*. Hoboken, NJ: Wiley-IEEE Press, 2004, ISBN: 0-471-47378-2.
- [8] P. M. Woodward, *Probability and Information Theory, with Applications to Radar*. London: Pergamon Press, 1953.
- [9] F. J. Harris, "On the use of windows for harmonic analysis with the discrete fourier transform," *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51–83, 1978. DOI: 10.1109/PROC.1978.10837
- [10] S. D. Blunt, M. Cook, J. Jakabosky, J. De Graaf, and E. Perrins, "Polyphase-coded FM (PCFM) radar waveforms, part I: Implementation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 3, pp. 2218–2229, Jul. 2014, ISSN: 0018-9251, 1557-9603, 2371-9877. DOI: 10.1109/TAES.2014.130361 Accessed: Oct. 2, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/6965769/>
- [11] J. Jakabosky, S. D. Blunt, and B. Himed, "Waveform design and receive processing for nonrecurrent nonlinear FMCW radar," in *2015 IEEE Radar Conference (RadarCon)*, May 2015, pp. 1376–1381. DOI: 10.1109/RADAR.2015.7131210 Accessed: Oct. 14, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/7131210/>
- [12] J. C. Harrington et al., "Challenges and Prospective Solutions for Non-Uniform Radar Waveforms in a Shared Spectrum," in *2024 IEEE Radar Conference (RadarConf24)*, May 2024, pp. 1–6. DOI: 10.1109/RadarConf2458775.2024.10548980 Accessed: Oct. 6, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/10548980/>
- [13] S. D. Blunt et al., "Principles and Applications of Random FM Radar Waveform Design," *IEEE Aerospace and Electronic Systems Magazine*, vol. 35, no. 10, pp. 20–28, Oct. 2020, ISSN: 1557-959X. DOI: 10.1109/MAES.2019.2953763 Accessed: Oct. 14, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/9216208/>
- [14] F. Gini, "Grand challenges in radar signal processing," *Frontiers in Signal Processing*, vol. 1, 2021. DOI: 10.3389/frsip.2021.664232

- [15] A. Aubry, A. De Maio, B. Jiang, and S. Zhang, "Ambiguity Function Shaping for Cognitive Radar Via Complex Quartic Optimization," *IEEE Transactions on Signal Processing*, vol. 61, no. 22, pp. 5603–5619, Nov. 2013, ISSN: 1941-0476. DOI: 10.1109/TSP.2013.2273885 Accessed: Oct. 29, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/6563125/>
- [16] A. Aubry, A. DeMaio, A. Farina, and M. Wicks, "Knowledge-Aided (Potentially Cognitive) Transmit Signal and Receive Filter Design in Signal-Dependent Clutter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 1, pp. 93–117, Jan. 2013, ISSN: 1557-9603. DOI: 10.1109/TAES.2013.6404093 Accessed: Sep. 24, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/6404093/>
- [17] A. Aubry et al., "Assessing Radar Waveforms for Spectral Coexistence via the PARSAX System," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 60, no. 5, pp. 6671–6684, Oct. 2024, ISSN: 0018-9251, 1557-9603, 2371-9877. DOI: 10.1109/TAES.2024.3412867 Accessed: Sep. 25, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/10555046/>
- [18] V. Carotenuto et al., "Experimental Evaluation of Radar Waveforms for Spectral Coexistence using the PARSAX radar," in *2023 IEEE 10th International Workshop on Metrology for AeroSpace (MetroAeroSpace)*, Jun. 2023, pp. 389–394. DOI: 10.1109/MetroAeroSpace57412.2023.10190037 Accessed: Oct. 6, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/10190037/>
- [19] A. Aubry, V. Carotenuto, A. De Maio, A. Farina, A. Izzo, and R. S. L. Moriello, "Assessing power amplifier impairments and digital predistortion on radar waveforms for spectral coexistence," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 1, pp. 635–650, 2021.
- [20] V. Carotenuto, A. Aubry, A. De Maio, A. Izzo, A. Farina, and R. S. L. Moriello, "Power amplifier distortions on radar signals for spectral coexistence," in *2021 Signal Processing Symposium (SPSymposium)*, 2021, pp. 35–39.
- [21] D. G. Felton, J. W. Owen, and S. D. Blunt, "Random FM Waveforms Jointly Optimized for Delay-Doppler Ambiguity Shaping," in *2025 IEEE International Radar Conference (RADAR)*, May 2025, pp. 1–6. DOI: 10.1109/RADAR52380.2025.11031843 Accessed: Oct. 1, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/11031843/>
- [22] J. Owen, C. Mohr, B. Ravenscroft, S. Blunt, B. Kirk, and A. Martone, "Real-Time Experimental Demonstration and Evaluation of Open-Air Sense-and-Notch Radar," in *2022 IEEE Radar Conference (RadarConf22)*, Mar. 2022, pp. 1–6. DOI: 10.1109/RadarConf2248738.2022.9764258 Accessed: Jan. 26, 2026. [Online]. Available: <https://ieeexplore.ieee.org/document/9764258/>
- [23] B. Ravenscroft, J. W. Owen, J. Jakabosky, S. D. Blunt, A. F. Martone, and K. D. Sherbondy, "Experimental demonstration and analysis of cognitive spectrum sensing and notching for radar," *IET Radar, Sonar & Navigation*, vol. 12, no. 12, pp. 1466–1475, 2018, ISSN: 1751-8792. DOI: 10.1049/iet-rsn.2018.5379 Accessed: Oct. 13, 2025. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1049/iet-rsn.2018.5379>
- [24] A. F. Martone et al., "Practical aspects of cognitive radar," in *2020 IEEE Radar Conference (RadarConf20)*, 2020, pp. 1–6. DOI: 10.1109/RadarConf2043947.2020.9266646
- [25] A. F. Martone, R. M. Buehrer, and D. M. McNamara, "Emerging Trends in Radar: Metacognitive Radar Networks for the Next Generation of Intelligent Sensing," *IEEE Aerospace and Electronic Systems Magazine*, vol. 40, no. 6, pp. 114–120, Jun. 2025, ISSN: 1557-959X. DOI: 10.1109/MAES.2025.3539607 Accessed: Sep. 24, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/10878308/>
- [26] V. Saarinen and V. Koivunen, "Radar Waveform Synthesis Using Generative Adversarial Networks," in *2020 IEEE Radar Conference (RadarConf20)*, Sep. 2020, pp. 1–6. DOI: 10.1109/RadarConf2043947.2020.9266709 Accessed: Oct. 13, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/9266709/>
- [27] M. Karásek, F. T. Muijres, C. De Wagter, B. D. Remes, and G. C. de Croon, "A tailless aerial robotic flapper reveals that flies use torque coupling in rapid banked turns," *Science*, vol. 361, no. 6407, 2018. DOI: 10.1126/science.aat0350

- [28] Ornadyne. “Reconnaissance redefined,” Accessed: May 21, 2026. [Online]. Available: <https://ornadyne.com/>
- [29] G. Cui, A. De Maio, A. Farina, and J. Li, Eds., *Radar Waveform Design Based on Optimization Theory* (Radar, Sonar and Navigation). SciTech Publishing, 2020.
- [30] A. Ziemann and C. A. Metzler, *Adaptive LPD radar waveform design with generative deep learning*, 2024. arXiv: 2403.12254 [eess.SP].
- [31] J. W. Owen et al., “Analysis of Spectrally Efficient Random FM Radar Waveforms,” in *2024 International Radar Conference (RADAR)*, Oct. 2024, pp. 1–6. DOI: 10.1109/RADAR58436.2024.10993723 Accessed: Oct. 14, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/10993723/>
- [32] E. R. Biehl, C. A. Mohr, B. Ravenscroft, and S. D. Blunt, “Assessment of Constant Envelope OFDM as a Class of Random FM Radar Waveforms,” in *2020 IEEE Radar Conference (RadarConf20)*, Sep. 2020, pp. 1–6. DOI: 10.1109/RadarConf2043947.2020.9266427 Accessed: Oct. 14, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/9266427/>
- [33] T. J. Kramer, E. R. Biehl, M. B. Heintzelman, S. D. Blunt, and E. D. Steinbach, “Compact Parameterization of Nonrepeating FMCW Radar Waveforms,” in *2023 IEEE Radar Conference (RadarConf23)*, May 2023, pp. 1–6. DOI: 10.1109/RadarConf2351548.2023.10149578 Accessed: Sep. 29, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/10149578/>
- [34] R. Sutton. “The Bitter Lesson,” Accessed: Nov. 7, 2025. [Online]. Available: <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>
- [35] A. L. Sènica, P. A. C. Marques, and M. A. T. Figueiredo, “Artificial intelligence applications in noise radar technology,” *IET Radar, Sonar & Navigation*, vol. 18, no. 7, pp. 986–1001, 2024. DOI: 10.1049/rsn2.12503
- [36] R. Palamà, F. Fioranelli, M. Ritchie, M. Inggs, S. Lewis, and H. Griffiths, “Measurements and discrimination of drones and birds with a multi-frequency multistatic radar system,” *IET Radar, Sonar & Navigation*, vol. 15, no. 8, pp. 841–852, 2021. DOI: 10.1049/rsn2.12060
- [37] R. Palamà et al., “First measurements with NeXtRAD, a polarimetric X/L-band radar network,” in *2017 IEEE Radar Conference (RadarConf)*, 2017. DOI: 10.1109/RADAR.2017.7944474
- [38] S. Z. Gurbuz, H. D. Griffiths, A. Charlish, M. Rangaswamy, M. S. Greco, and K. Bell, “An Overview of Cognitive Radar: Past, Present, and Future,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 34, no. 12, pp. 6–18, Dec. 2019, ISSN: 1557-959X. DOI: 10.1109/MAES.2019.2953762 Accessed: Oct. 13, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/8961364/>
- [39] S. Haykin, “Cognitive radar: A way of the future,” *IEEE Signal Processing Magazine*, vol. 23, no. 1, pp. 30–40, Jan. 2006, ISSN: 1558-0792. DOI: 10.1109/MSP.2006.1593335 Accessed: Oct. 29, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/1593335/>
- [40] E. Elelimy, D. Szepesvari, M. White, and M. Bowling, “Rethinking the foundations for continual reinforcement learning,” 2025. arXiv: 2504.08161. [Online]. Available: <https://arxiv.org/abs/2504.08161>
- [41] K. Javed and R. S. Sutton, “The Big World Hypothesis and its Ramifications for Artificial Intelligence,” in *Finding the Frame: An RLC Workshop for Examining Conceptual Frameworks*, 2024. [Online]. Available: <https://openreview.net/forum?id=Sv7DazuCn8>
- [42] Y. Wu, G. D. Singh, M. Beikmirza, L. C. N. de Vreede, M. Alavi, and C. Gao, “OpenDPD: An Open-Source End-to-End Learning & Benchmarking Framework for Wideband Power Amplifier Modeling and Digital Pre-Distortion,” in *2024 IEEE International Symposium on Circuits and Systems (ISCAS)*, May 19, 2024, pp. 1–5. DOI: 10.1109/ISCAS58744.2024.10558162 arXiv: 2401.08318 [cs]. Accessed: Oct. 31, 2025. [Online]. Available: <http://arxiv.org/abs/2401.08318>
- [43] D. Eustice, C. Baylis, L. Cohen, and R. J. Marks, “Effects of power amplifier nonlinearities on the radar ambiguity function,” in *2015 IEEE Radar Conference (RadarCon)*, 2015, pp. 1725–1729. DOI: 10.1109/RADAR.2015.7131266

- [44] J. E. Quirk, M. B. Heintzelman, D. B. Herr, P. M. McCormick, S. D. Blunt, and B. Maio, "Aspect Angle Classification via Physically Realizable Matched Illumination Waveform Design," in *2025 IEEE International Radar Conference (RADAR)*, May 2025, pp. 1–6. DOI: 10.1109/RADAR52380.2025.11031675 Accessed: Oct. 1, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/11031675/>
- [45] B. Correll, T. D. Bufler, C. N. Swanson, and R. M. Narayanan, "A formal study of the doppler tolerance of costas and sudoku waveforms," in *2021 IEEE Radar Conference (RadarConf21)*, IEEE, 2021, pp. 1–6.
- [46] J. Fraka, T. Higgins, and J. Owen, "Real-Time Waveform-Diverse Pulse-Doppler Demo via Microwave Radar-in-a-Briefcase (MicRIB)," in *2024 IEEE Radar Conference (RadarConf24)*, May 2024, pp. 1–6. DOI: 10.1109/RadarConf2458775.2024.10548690 Accessed: Sep. 19, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/10548690/>
- [47] S. Stephany, B. Schweizer, C. Knill, and C. Waldschmidt, "Impact of an Automotive Chirp-Sequence Interferer on a Wideband Pseudo-Noise Radar," in *2019 International Conference on Electromagnetics in Advanced Applications (ICEAA)*, Sep. 2019, pp. 0859–0862. DOI: 10.1109/ICEAA.2019.8879023 Accessed: Sep. 25, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/8879023/>
- [48] G. Zook, P. M. McCormick, S. D. Blunt, C. Allen, and J. Jakabosky, "Dual-polarized FM noise radar," in *International Conference on Radar Systems (Radar 2017)*, Oct. 2017, pp. 1–5. DOI: 10.1049/cp.2017.0442 Accessed: Oct. 14, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/8367527/>
- [49] R. Gourova, R. Pribić, and A. Yarovoy, "Theory and practice of an Alltop waveform," in *2016 17th International Radar Symposium (IRS)*, May 2016, pp. 1–6. DOI: 10.1109/IRS.2016.7497324 Accessed: Nov. 11, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/7497324/>
- [50] L. Silva, A. Rathnayake, H. Rezaei, and N. Rajatheva, "Transformer neural network-based behavioral modeling and predistortion for wideband power amplifiers," in *2024 IEEE 35th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2024, pp. 1–6. DOI: 10.1109/PIMRC59610.2024.10817328
- [51] A. Vaswani et al., "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, Curran Associates, Inc., 2017.
- [52] H. Zou et al., *RF-GPT: Teaching AI to see the wireless world*, 2026. arXiv: 2602.14833 [eess.SP].
- [53] D. Hafner, T. Lillicrap, M. Norouzi, and J. Ba, "Mastering diverse domains through world models," *arXiv preprint arXiv:2301.04104*, 2023. [Online]. Available: <https://arxiv.org/abs/2301.04104>
- [54] M. A. Ritchie, F. Fioranelli, H. Borrión, and H. D. Griffiths, "Multistatic micro-Doppler radar feature extraction for classification of unloaded/loaded micro-drones," *IET Radar, Sonar & Navigation*, vol. 11, no. 1, pp. 116–124, 2017. DOI: 10.1049/iet-rsn.2016.0063
- [55] N. Kruse, R. Guendel, F. Fioranelli, and A. Yarovoy, "Distributed radar fusion for extended target location and velocity reconstruction," in *2024 IEEE Radar Conference (RadarConf24)*, Denver, USA, 2024. DOI: 10.1109/RadarConf2458775.2024.10548772
- [56] L. Yan, S. Han, C. Hao, D. Orlando, and G. Ricci, "Innovative cognitive approaches for joint radar clutter classification and multiple target detection in heterogeneous environments," *IEEE Transactions on Signal Processing*, vol. 71, pp. 1010–1022, 2023. DOI: 10.1109/TSP.2023.3250084
- [57] C. Rapp, "Effects of HPA-nonlinearity on a 4-DPSK/OFDM signal for the mobile satellite channel," in *Proceedings of the 2nd European Conference on Satellite Communications (ECSC-3)*, vol. 332, 1991, pp. 179–184.
- [58] A. Hebb, A. Yarovyi, and F. Fioranelli, "Learned pre-distortion of diverse radar waveforms," Work in progress. Draft submitted to EuRAD 2026, 2026.
- [59] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention*, Springer, 2015, pp. 234–241.

- [60] M. Jaderberg et al., "Population based training of neural networks," 2017. arXiv: 1711.09846.
- [61] T. Epailly and O. Krasnov, "Implementation of the PARSAX Polarimetric Radar Real-Time Sampling and Signal Processing Methods on a Vadatech-Xilinx FPGA board," Oct. 7, 2024.
- [62] D. Hafner, W. Yan, and T. Lillicrap, "Training agents inside of scalable world models," *arXiv preprint arXiv:2509.24527*, 2025. [Online]. Available: <https://arxiv.org/abs/2509.24527>
- [63] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017. arXiv: 1707.06347.
- [64] N. A. Goodman, P. R. Venkata, and M. A. Neifeld, "Adaptive waveform design and sequential hypothesis testing for target recognition with active sensors," *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, no. 1, pp. 105–113, 2007.
- [65] D. B. Herr, P. S. Raju, and J. M. Stiles, "Information theoretic waveform design with applications to adaptive-on-transmit radar," *IET Radar, Sonar & Navigation*, vol. 18, no. 1, pp. 222–234, 2024, ISSN: 1751-8792. DOI: 10.1049/rsn2.12478 Accessed: Oct. 2, 2025. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1049/rsn2.12478>
- [66] E. P. Lys and Z. Li, "PARSAX: Radar Specifications," Feb. 2007.
- [67] I. Tektronix, *AWG5000 and AWG7000 Series Arbitrary Waveform Generators Programmer Manual*, Tektronix, Inc., Sep. 22, 2011.
- [68] H. V. Sethuraman, A. Yarovoy, and F. Fioranelli, "Classification of unmanned aerial vehicles carrying payloads with polarimetric radar," in *2021 18th European Radar Conference (EuRAD)*, 2022, pp. 365–368. DOI: 10.23919/EuRAD50154.2022.9784557

# A

## Learned Pre-Distortion of Diverse Radar Waveforms

We submitted this paper to EuRAD 2026, which accepted it as an oral presentation tentatively scheduled for 8 October 2026.

# Learned Pre-Distortion of Diverse Radar Waveforms

Amos Hebb, Alexander Yarovoy, Francesco Fioranelli  
MS3 Group, Delft University of Technology, Netherlands  
a.n.hebb@tudelft.nl, a.yarovoy@tudelft.nl, f.fioranelli@tudelft.nl

**Abstract**—In this paper, the problem of distortion of radar waveforms caused by power amplifiers is investigated. Existing digital pre-distortion techniques address this for a single waveform, but must be re-calibrated whenever the waveform type changes. We present a single network that pre-distorts 96 radar-related waveforms across 3 families (linear-frequency-modulated chirps, smoothed phase-coded waveforms, and orthogonal frequency-division multiplexed waveforms) with no waveform-specific parameters or explicit waveform labels. Training places the power amplifier in the optimization loop, requiring no separate identification step or amplifier surrogate. At the saturation knee, the model reduces the root-mean-square distortion between reference and post-amplification waveforms in IQ by 67–75% across families of waveforms, and generalizes to families held out entirely during training.

**Keywords**—digital pre-distortion, power amplifier, waveform diversity, radar, deep learning

## I. INTRODUCTION

Modern radar systems increasingly rely on waveform agility to manage spectral coexistence, improve target discrimination, and implement cognitive sensing strategies [1]. Switching between diverse transmit waveforms on a pulse-to-pulse or burst-to-burst basis offers clear advantages, but it also complicates the transmit chain. Power amplifiers are inherently nonlinear devices, compressing and distorting the waveform when driven near saturation to maximize efficiency. For a single fixed waveform, one can characterize this distortion once and compensate with a static look-up table or polynomial pre-distortion filter. However, per-waveform re-calibration is not only expensive during manufacture of the radar system, but, for an operational system that transmits a library of waveforms with different bandwidths, sweep directions, or modulation schemes, the distortion can change from pulse to pulse and a static correction no longer suffices.

The radar community traditionally handles this by reducing the amplifier gain (sacrificing range) or restricting to constant-envelope waveforms [1], but both approaches reduce performance. Even for a waveform designed to be constant-envelope, Kumbul *et al.* [2] show that the phase lag compensation required for practical de-chirped reception introduces amplitude variation that drives the amplifier into its nonlinear region. The pre-distortion problem and the waveform diversity problem are therefore compounded and not separable.

Polynomial pre-distortion has been applied to radar waveforms by Aubry *et al.* [3] and Herndon and Yearly [4], but the coefficients are tied to a specific waveform and require unit specific re-calibration when this changes. Learned




Figure 1. Block diagram of the signal chain to describe the proposed approach. IQWaveNet maps a desired complex-valued waveform in baseband  $x_{\text{ref}}$  to a pre-distorted command waveform  $x_{\text{cmd}}$ . The hardware stages upsample this to a real-valued intermediate-frequency rate signal, that gets amplified by the PA and decimated back (specifically  $8\times$  in the case the TU Delft PARSAX radar considered in this work as an example). Dashed line shows the bypass path without pre-distortion. We call the signal at the output of the proposed full chain  $\text{PA}(\cdot)$  in the rest of this work.

pre-distortion is active in communications [5], [6] but targets quasi-Gaussian modulations, not the more structured, deterministic waveforms used in radar.

In this paper, we propose a WaveNet-style [7] dilated convolutional network trained with the power amplifier in the loop, which takes a desired waveform in and outputs its suitable pre-distorted version directly. The proposed network operates on the raw in-phase/quadrature (IQ) samples and can handle different waveform classes such as Linear Frequency Modulated (LFM) and Phase-Coded Frequency Modulated Continuous Wave (PC-FMCW) chirps or Orthogonal Frequency Division Multiplexing (OFDM), without explicit, waveform-specific parameters.

The contributions of our work are:

- 1) A single pre-distortion model for 96 radar waveforms across 3 families (LFM, PC-FMCW, OFDM waveforms), trained end-to-end through a differentiable power amplifier distortion model with thermal memory.
- 2) Correction of both amplitude and phase distortion, including time-varying thermal drift that no static polynomial filter can follow.
- 3) Demonstration of generalization to entire waveform families not seen during training.

The rest of the paper is organized as follows. Section II introduces the system model of the proposed approach, with details of the pre-distortion network in Section III. Section IV presents the initial results with conclusions in Section V.

## II. SYSTEM MODEL

As an example of an operational radar to present the proposed approach, we model the transmit chain of the PARSAX research radar at Delft University of Technology [3]. Figure 1 shows the block diagram of the signal path, with the overall goal of finding the pre-distorted signal  $x_{\text{cmd}}$  such that  $\text{PA}(x_{\text{cmd}}) \approx x_{\text{ref}}$ . This essentially means that the pre-distorted

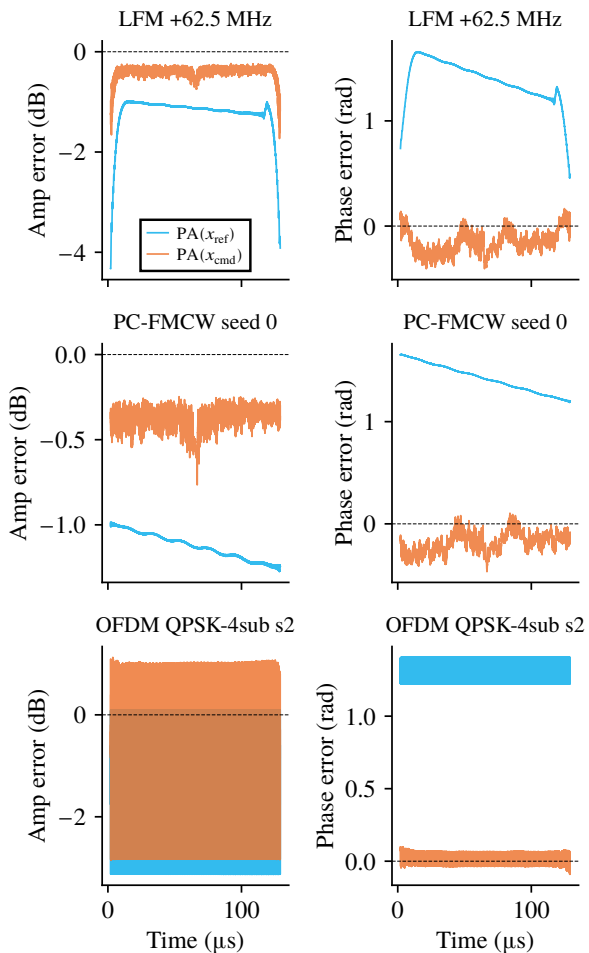


Figure 2. Amplitude and phase error across waveforms for 1 example per waveform family. The network output (orange) tracks the desired waveform more closely than the uncorrected baseline (blue). Error is roughly flat over time, showing that the proposed pre-distortion network captures thermal drift.

signal passed through the Power Amplifier (PA) is as similar as possible to the desired, reference waveform.

Since the intention is to move the proposed approach to hardware in the loop, here we make no attempt to faithfully reproduce any particular amplifier model. To keep the effect of the PA general, we combine a Rapp [8] static nonlinearity with a first-order thermal memory that makes the distortion time-varying within a single waveform. The output sample at time  $t$  takes the form:

$$y(t) = x(t) \cdot G_{\text{Rapp}}(|x(t)|) \cdot \max(0, 1 - \alpha s(t)) \cdot e^{j\beta s(t)} \quad (1)$$

where  $G_{\text{Rapp}}(r) = (1 + (r/v_{\text{sat}})^{2p})^{-1/(2p)}$  with  $p = 2$ , and  $s(t)$  is a first-order causal filter on instantaneous input power with a time constant of approximately  $100 \mu\text{s}$  at  $62.5 \text{ MS/s}$ . The thermal coefficients ( $\alpha = 2.0$ ,  $\beta = 5.0 \text{ rad}$ ) produce roughly 1 dB of amplitude taper and  $10^\circ$  of phase drift across a  $131 \mu\text{s}$  waveform that memoryless models struggle to capture.

As an example, blue curves in Figure 2 show amplitude error ramps and phase drifts over time caused by the aforementioned PA distortion model, both signatures of

the thermal state accumulating. Notably, we apply the nonlinearity at the intermediate-frequency rate signal to produce in-band distortion and harmonics which, roughly, resemble experimentally collected data. We set  $v_{\text{sat}} = 0.65$  with a waveform amplitude scale of 0.80, placing the peak at 125% of  $v_{\text{sat}}$ , right at the saturation knee.

#### A. Waveform Library

To evaluate the proposed approach, we train a single neural network on 96 waveforms drawn from 3 distinct families.

a) *Linear-frequency-modulated chirps*: Bandwidths are linearly spaced from  $-62.5$  to  $+62.5$  MHz (negative means downchirp). All the waveforms are simulated at the 125 MHz intermediate-frequency center, and present the amplifier with a nearly constant envelope where only instantaneous frequency varies across the waveform.

b) *Smoothed phase-coded waveforms*: Per Kumbul *et al.* [2], PC-FMCW waveforms are modeled as a 50 MHz upchirp carrier modulated by 16 random binary chips with Gaussian-smoothed transitions, with diversity coming from 32 different random seeds. Phase lag compensation, applied in the frequency domain before transmission so that the de-chirping receiver sees clean range sidelobes, introduces amplitude variation that drives the amplifier envelope into its nonlinear region. Even for a waveform designed to be constant-envelope, the processing required for its practical use makes it not.

c) *OFDM waveforms*: These specifically include 4 subcarrier counts (4, 8, 16, 32) with 2 modulations (binary and quadrature phase-shift keying) and 4 random seeds each. Increasing subcarrier count raises the peak-to-average power ratio, creating progressively more severe peak clipping and a wider range of distortion regimes for the network to learn. These waveforms present the hardest challenge. The amplifier sees a wide amplitude distribution rather than the nearly constant envelope of the other 2 families.

### III. LEARNED PRE-DISTORTION

This section presents the architecture of the neural network implementing the proposed pre-distortion.

#### A. Architecture: Residual IQWaveNet

IQWaveNet is a WaveNet-style [7] dilated convolutional network adapted from audio synthesis. The key change to adapt to radar signals is adding a second input and output channel to explicitly represent phase. The network takes 2 channels (IQ components of  $x_{\text{ref}}$ ) as input, and produces 2 output channels (the pre-distortion correction).

With respect to architecture, 9 dilated residual blocks with kernel size 15 and dilation factors 1, 2, 4, ..., 256 give a receptive field of 14 309 samples, covering the full 8 192 sample waveform. Each block applies SiLU “swish” activation and group normalisation. SiLU is a smooth, non-zero-gradient alternative to ReLU, popular in WaveNet and audio synthesis architectures. A bottleneck self-attention layer subsamples by a factor of 32, applies multi-head attention on the resulting 256 tokens, and broadcasts back to the full sequence, giving

the model a global communication channel that dilated convolutions can only approximate.

The network operates entirely on the raw IQ samples and must infer waveform structure from the signal itself. Earlier experiments used a class embedding to condition the network on the waveform type. However, removing this did not impact performance, confirming the model infers waveform structure from the signal itself.

The full model computes the following relation:

$$x_{\text{cmd}} = x_{\text{ref}} + s \cdot f_{\theta}(x_{\text{ref}}) \quad (2)$$

where  $f_{\theta}$  represents the IQWaveNet network, and  $s$  is a learned scalar initialised to 0.1. The network only learns the correction, not the full waveform. The scale  $s$  converges to approximately 0.2, confirming that the correction remains a modest perturbation. The model  $f_{\theta}$ , like any WaveNet network architecture, first projects the input  $x_{\text{ref}}$  into a number of “channels”. This internal dimension, about 5M parameters at our selected 128 channels, enables the network to model complex non-linearities in the waveform. The final layer of  $f_{\theta}$  then projects these features back into the 2-channel IQ to produce  $x_{\text{cmd}}$ .

### B. Training

The key idea is power amplifier in the loop training. The loss function passes the network’s output through the full simulated signal chain shown in Figure 1 and measures how well the amplifier output matches the desired waveform, requiring no separate identification step.

The training loss for a batch of waveforms is:

$$\mathcal{L} = \frac{1}{B} \sum_{i=1}^B \left\| \text{PA}(f_{\theta}(x_{\text{ref}}^{(i)})) - x_{\text{ref}}^{(i)} \right\|^2 \quad (3)$$

where  $\text{PA}(\cdot)$  denotes the full chain of upsampling, quantisation, amplifier, and downsampling described in Section II.

With the 96 waveforms across 3 families, each training step presents the entire library in a single batch. We train using Adam with a learning rate of  $10^{-4}$ , 200 warmup steps, cosine decay to 1/30 of peak, and global gradient-norm clipping at 1.0. An exponential moving average of the parameters with decay 0.9999 provides the final model.

Fully implemented in JAX, we differentiate through the amplifier model used in the simulations of this paper, but the loss requires only input-output pairs, so that the same procedure can also work with gradient-free methods applied to a physical amplifier. Training 65 000 steps, 96 classes, full width batches, 128 channels, takes approximately half an hour on a single NVIDIA GeForce RTX 5090.

## IV. RESULTS

### A. Pre-Distortion Performance

Table 1 summarizes improvement in RMS when using the network trained on all 96 waveforms. We compare the proposed pre-distortion to passing the reference waveform directly in as a command waveform (dashed line in Figure 1, as currently done with the TU Delft PARSAX radar). Baseline

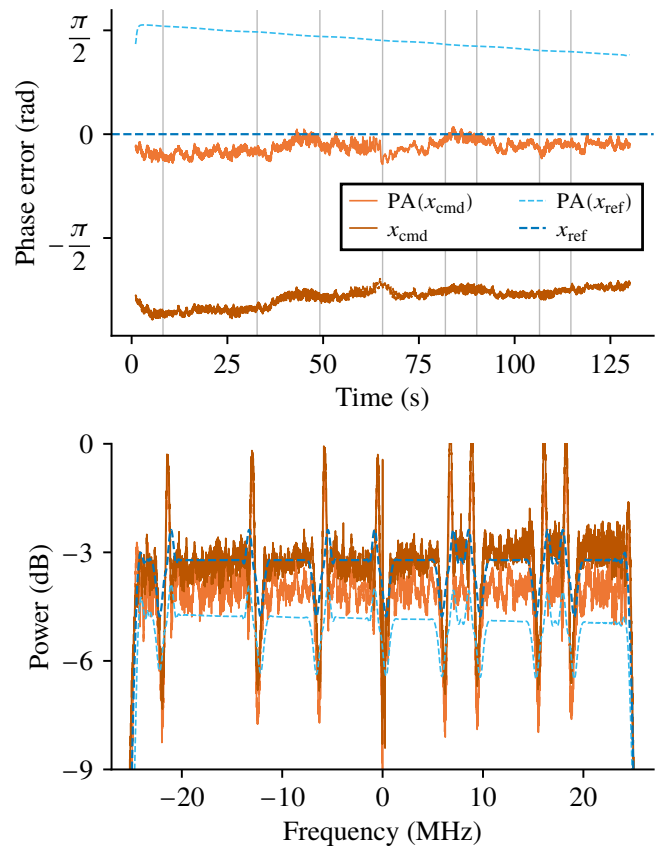


Figure 3. Smoothed PC-FMCW waveform [2]. Dashed lines are no IQWaveNet applied. Darker blue and brown are waveforms passed into amplifier model, cyan and orange lighter colors for post amplifier chain outputs. Top: de-chirped phase, where the network output  $\text{PA}(x_{\text{cmd}})$  nearly covers the desired  $x_{\text{ref}}$ . Bottom: zoomed-in power spectrum where again  $\text{PA}(x_{\text{cmd}})$  is closer to  $x_{\text{ref}}$  than  $\text{PA}(x_{\text{ref}})$  showing improvement.

RMS is the mean  $\text{RMS}(\text{PA}(x_{\text{ref}}))$ . Network RMS is the mean  $\text{RMS}(\text{PA}(x_{\text{cmd}}))$ . Improvement shows the range of values for  $1 - \text{RMS}(\text{PA}(x_{\text{cmd}}))/\text{RMS}(\text{PA}(x_{\text{ref}}))$  across all 32 examples from a waveform family, expressed as the nearest percentage.

Table 1. Pre-distortion performance by waveform family

Family	Baseline RMS	Network RMS	Improvement
LFM	0.154	0.042	67–75 %
PC-FMCW	0.146	0.039	72 %
OFDM	0.161	0.098	11–61 %

Outputs of the proposed approach  $x_{\text{cmd}}$  in Figure 2 show the two biggest contributors to loss, amplitude error and phase error, across the waveform duration for 1 example from each family. The model learns phase corrections across the full waveform duration, and the error is roughly flat over time, confirming that the model captures the thermal drift rather than only the static nonlinearity. The model struggles more with amplitude than phase, but the amplitude error it does produce lacks the systematic ramp that dominates the uncorrected baseline.

Smoothed PC-FMCW waveforms [2] are designed with




Figure 4. OFDM constellation. The PA distortion effect pushes blue symbols away from the dark-blue reference rings. After usage of the proposed approach, the pre-distorted brown symbols land closer to the reference (orange).

Table 2. Generalization with phase-coded waveforms held out.

Family	Improvement	Status
LFM	68 %	trained
PC-FMCW	51 %	<b>held out</b>
OFDM	26 %	trained

constant envelope in mind, but the phase lag compensation required for de-chirped reception introduces amplitude variation. Figure 3 shows the model recovering de-chirped phase to within a few degrees across the full waveform. Notably, the  $x_{\text{cmd}}$  signal is far ahead of the  $x_{\text{ref}}$  signal, showing that the pre-distorter learns to anticipate the phase delay the PA introduces.

OFDM waveforms are the hardest, since the amplifier sees everything from near-zero to deep saturation within a single waveform. Performance scales inversely with peak-to-average power ratio, but the constellation diagram (Figure 4) shows the network pulling received symbols toward their intended, ideal positions for small symbol sets, proving the value of the proposed approach.

### B. Generalization to Unseen Waveform Families

A waveform-diverse radar may eventually need to transmit waveforms not in the original training set. To test generalization, we train another version of the network on only the chirp and multi-carrier families (64 waveforms), holding out all 32 PC-FMCW waveforms entirely.

The held-out PC-FMCW family generalizes the chirps with phase transitions (Table 2), with notably small variance across all 32 random seeds (50.3–51.3 %). Notably, since the network receives no class information, this transfer comes from learning the amplifier’s transfer function itself, not from

memorizing per-waveform corrections. The overall numbers are lower than Table 1 due to high loss in OFDM overwhelming the shared residual scale parameter. Including the phase-coded waveforms during training raises their improvement from 51 % to 72 % (Table 1), so while the model generalizes well to unseen families, there remains a clear benefit to training on all families the system will transmit.

## V. CONCLUSION

We present a network trained to pre-distort radar waveforms with power amplifier in the optimization loop. Results show that a single network with no waveform-specific parameters can handle 96 waveforms across 3 families, correcting both static nonlinearity and time-varying thermal drift, and can generalize to families not seen during training. Real solid-state amplifiers exhibit higher-order and longer duration memory effects not yet present in this simulation, though the approach is designed with handling whole waveform intervals in mind. In future work, the simulation is planned to be replaced by using the PARSAX radar loopback measurement path, which provides exactly the input-output pairs the training loss of this paper requires. The broader goal is a system that jointly generates novel radar waveforms and pre-distorts them in a single forward pass, enabling waveform diversity without requiring amplifier linearity.

## ACKNOWLEDGMENT

This work is in part supported by Thales Germany and Holland High Tech within the WAVERAD project.

## REFERENCES

- [1] S. D. Blunt and E. L. Mokole, “Overview of radar waveform diversity,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 31, no. 11, Nov. 2016.
- [2] U. Kumbul, N. Petrov, C. S. Vaucher, and A. Yarvoy, “Smoothed phase-coded FMCW: Waveform properties and transceiver architecture,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 59, no. 2, pp. 1720–1733, Apr. 2023.
- [3] A. Aubry, V. Carotenuto, A. De Maio, F. Fioranelli, O. Krasnov, A. Yarvoy, and F. van der Zwan, “Assessing radar waveforms for spectral coexistence via the PARSAX system,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 60, no. 5, pp. 6671–6684, 2024.
- [4] M. Herndon and M. Yeary, “Real-time FPGA-based digital predistortion for improved amplifier performance in next generation phased arrays,” in *2022 IEEE Radar Conference (RadarConf22)*, 2022.
- [5] Y. Wu, G. D. Singh, M. Beikmirza, L. C. N. de Vreede, M. Alavi, and C. Gao, “OpenDPD: An open-source end-to-end learning and benchmarking framework for wideband power amplifier modeling and digital pre-distortion,” in *2024 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, May 2024, pp. 1–5.
- [6] C. Spano, D. Badini, L. Cazzella, and M. Matteucci, “Digital pre-distortion with deep reinforcement learning for 5G power amplifiers,” in *Proceedings of the 54th European Microwave Conference*, Paris, France, Sep. 2024.
- [7] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “WaveNet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
- [8] C. Rapp, “Effects of HPA-nonlinearity on a 4-DPSK/OFDM-signal for a digital sound broadcasting system,” in *Proceedings of the Second European Conference on Satellite Communications (ECSC-2)*, Liège, Belgium, Oct. 1991, pp. 179–184.