A Machine-Learning Approach to Analog/RF Circuit Testing*

Yiorgos Makris

TRELA

Testable and RELiable Architectures

Departments of Electrical Engineering & Computer Science

YALE UNIVERSITY

* Joint work with Dr. Haralampos Stratigopoulos (TIMA Lab, Grenoble, France), Prof. Petros Drineas (RPI) and Dr. Mustapha Slamani (IBM).
Partially funded by NSF, SRC/IBM, TI, and the European Commission.

Test and Reliability @ YALE

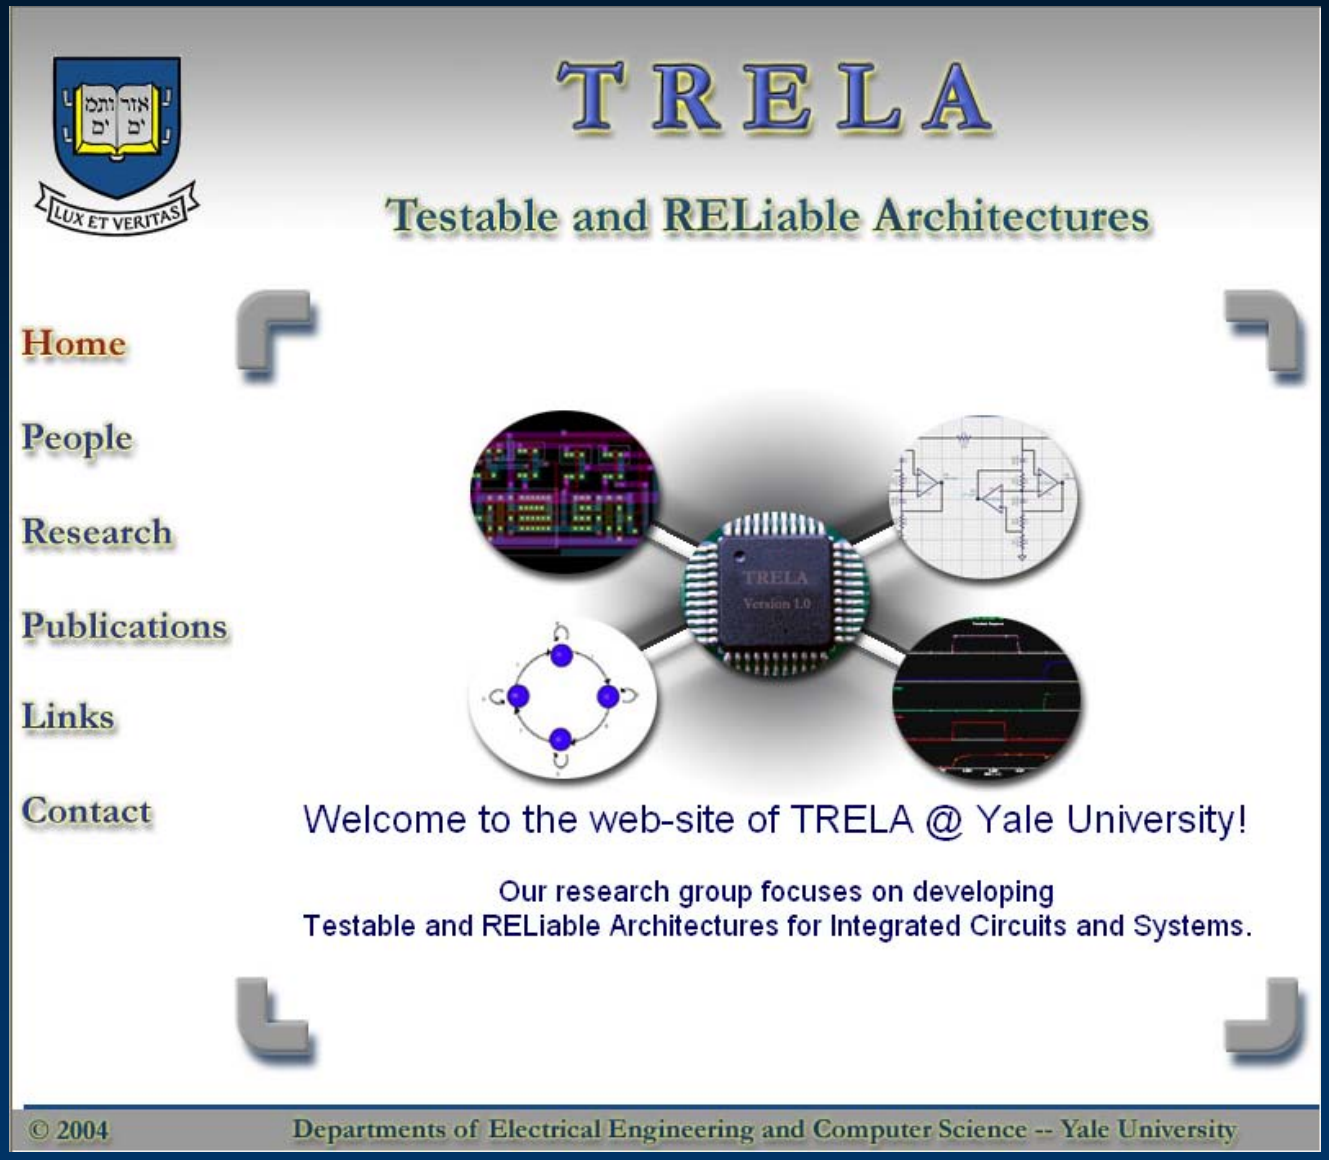

http://eng.yale.edu/trela

Research Areas

- Analog/RF circuits
  - Machine learning-based testing
  - Correlation mining for post-fabrication tuning and yield enhancement
  - Design of on-chip checkers and on-line test methods
  - Hardware Trojan detection in wireless cryptographic circuits
- Digital Circuits
  - Workload-driven error impact analysis in modern microprocessors
  - Logic transformations for improved soft-error immunity
  - Concurrent error detection / correction methods for FSMs
- Asynchronous circuits
  - Fault simulation & test generation for speed-independent circuits
  - Test methods for high-speed pipelines (e.g. Mousetrap)
  - Error detection and soft-error mitigation in burst-mode controllers

Presentation Outline

- Testing analog/RF circuits
- Machine learning-based testing
- Testing via a non-linear neural classifier
  - Construction and training
  - Selection of measurements
  - Test quality vs. test time trade-off
- Experimental results
- Analog/RF specification test compaction
- Stand-alone Built-in Self-Test (BIST)
- Performance calibration via post-fabrication tuning
- Yield enhancement
- Summary

Definition of Analog/RF Functionality

Symbol

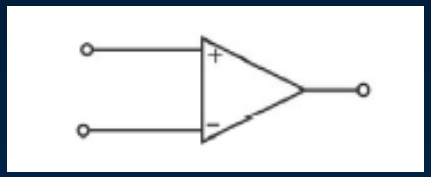

Transistor-Level

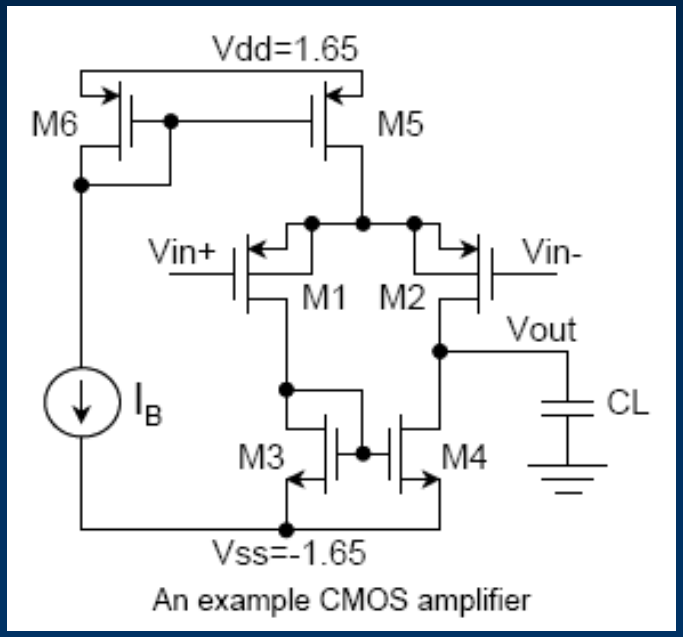

An example CMOS amplifier

Specifications

2.7V Electrical Characteristics

V^+ = 2.7V, V^- = 0V, V_{CM} = 1.35V, T_A = 25°C unless otherwise stated. Boldface limits apply over the Temperature Range.

Symbol	Parameter	Condition	Typ (Note 5)	Limit (Note 2)	Units
V_{OS}	Input Offset Voltage		0.05	1.0 1.5	mV max
V_{CM}	Input common-Mode Voltage Range	For CMRR ≥ 50dB		0	V min
			1.4	1.3	V max
CMRR	Common Mode Rejection Ratio	0V < V_{CM} < 1.3V	100	85 70	dB min
PSRR	Power Supply Rejection Ratio	V^+ = 2.7V to 5.0V	107	85 70	dB min
I_S	Supply Current		0.5	0.8 0.85	mA max
I_{IN}	Input Current		1.5	100	pA max
I_{OS}	Input Offset Current		0.2		pA
A_{VOL}	Voltage Gain	R_L = 10k Connect to $V^+/2$ V_O = 0.2V to 2.2V	120	110 95	dB min
		R_L = 2k Connect to $V^+/2$ V_O = 0.2V to 2.2V	120	100 85	
V_O	Positive Voltage Swing	R_L = 10k Connect to $V^+/2$	2.62	2.54 2.52	V min
		R_L = 2k Connect to $V^+/2$	2.62	2.54 2.52	
V_O	Negative Voltage Swing	R_L = 10k Connect to $V^+/2$	78	140 160	mV max
		R_L = 2k Connect to $V^+/2$	78	160 180	
I_O	Output Current	Sourcing, V_O = 0V V_{IN} (diff) = ±0.5V	12	6.0 1.5	mA min
		Sinking, V_O = 2.7V V_{IN} (diff) = ±0.5V	11	6.0 1.5	
e_n (10Hz)	Input Referred Voltage Noise		15.5		$nV/\sqrt{Hz}$
e_n (1kHz)	Input Referred Voltage Noise		7		$nV/\sqrt{Hz}$

Analog/RF IC Testing - Problem Definition

Comparison

Pre-layout or post-layout performances

Simulation

Actual silicon performances

Measurement

Characteristic	Specification Goal
Frequency	5150 to 5875 MHz
Bandwidth	>725 MHz
Gain	>15 dB
Gain Ripple	±0.5 dB
Noise Figure	<5 dB, 3 dB goal
Input IP3	>+5 dBm
VSWR, 50 Ohm	<1.5:1 (14 dB) input & output
Supply Voltage	+5 Volts only, goal
Size	60 X 60 mil ANACHIP

Design

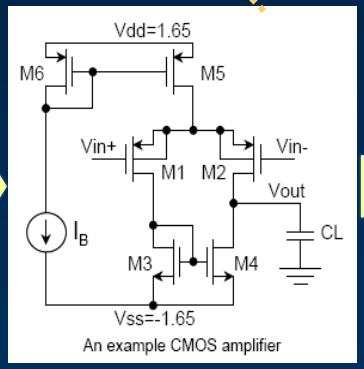

An example CMOS amplifier

Layout

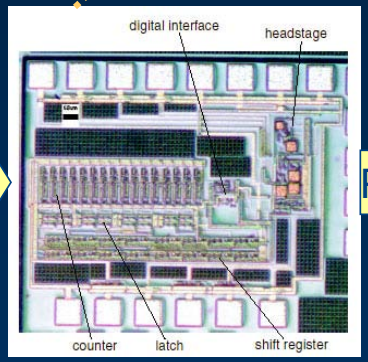

Fabrication

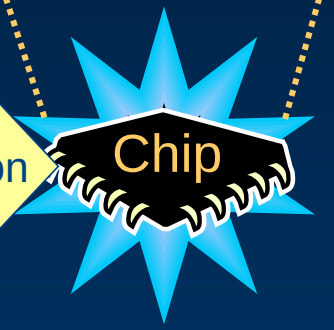

Verification:

- Targets design errors
- Once per design

Testing:

- Targets manufacturing defects
- Once per chip

Analog/RF IC Test - Industrial Practice

- Post-silicon production flow

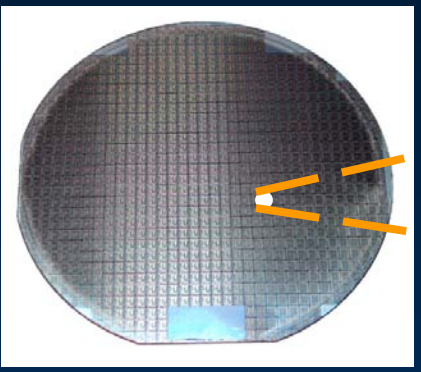

Wafer

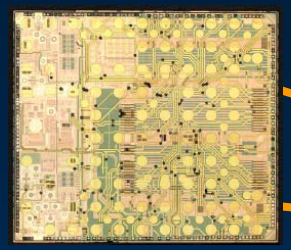

Die

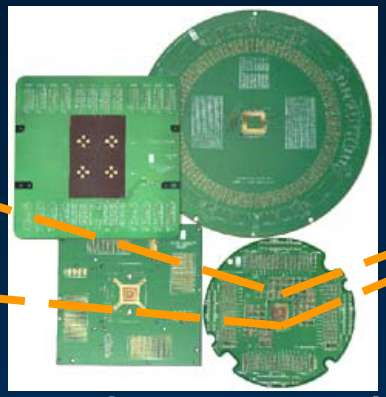

Interface Board

Automatic Test Equipment (ATE)

- Current practice is specification testing

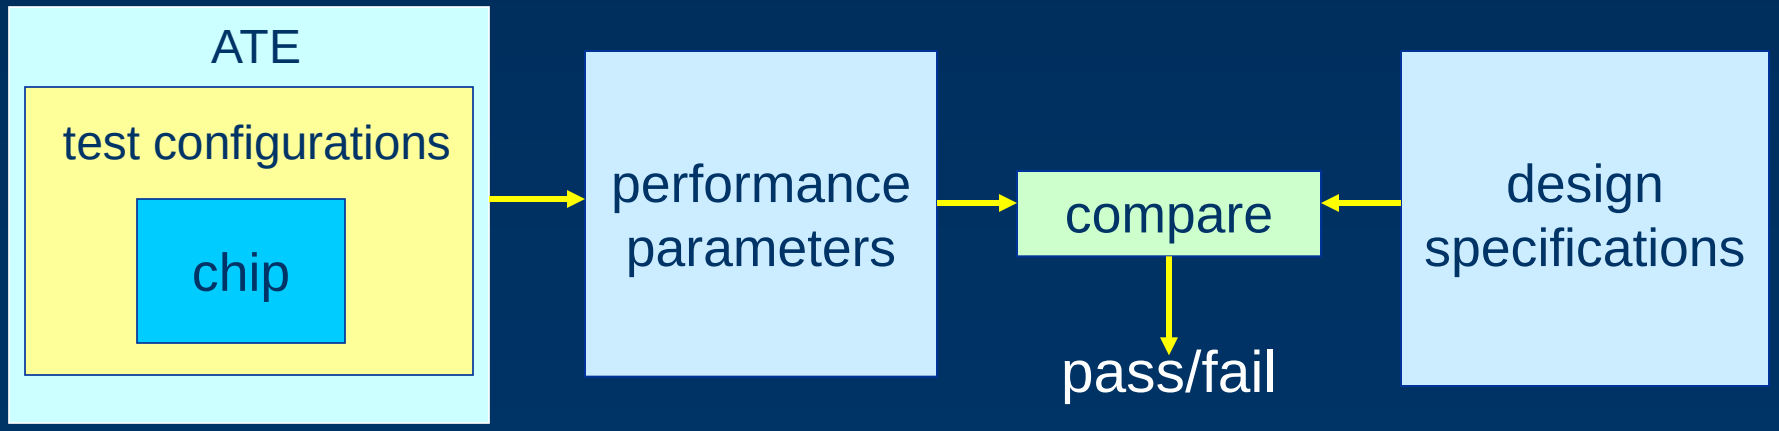

Limitations

Test Cost:

- Expensive ATE (multi-million dollar equipment)
- Specialized circuitry for stimuli generation and response measurement

Test Time:

- Multiple measurements and test configurations
- Switching and settling times

Alternatives?

- Fault-model based test – Never really caught on
- Machine learning-based (a.k.a. “alternate”) testing
  - Regression (*Variyam et al., TCAD’02*)
  - Classification (*Pan et al., TCAS-II’99, Lindermeir et al., TCAD’99*)

Machine Learning-Based Testing

General idea:

- Determine if a chip meets its specifications without explicitly computing the performance parameters and without assuming a prescribed fault model

How does it work?

- Infer whether the specifications are violated through a few simpler/cheaper measurements and information that is "learned" from a set of fully tested chips

Underlying assumption:

- Since chips are produced by the same manufacturing process, the relation between measurements and performance parameters can be statistically learned

Regression vs. Classification

Problem Definition:

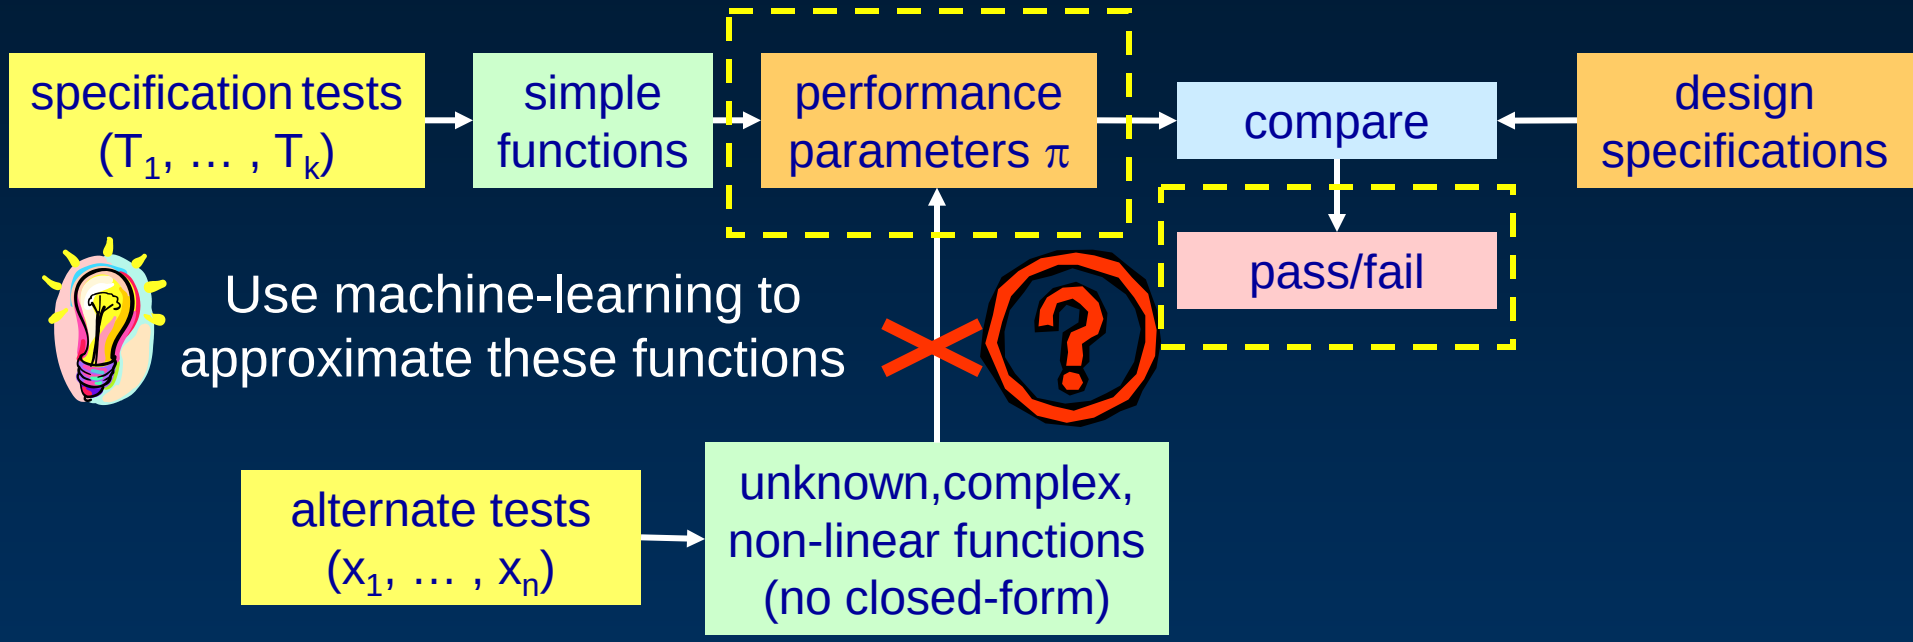

Regression:

- Explicitly learn these functions (i.e. approximate f:x → π)

Classification:

- Implicitly learn these functions (i.e. approximate f:x→Y,Y={pass/fail})

Overview of Classification Approach

training set of chips

simpler measurements

acquisition of measurement patterns

specification tests

pass/fail labels

projection on measurement space

LEARNING

learn boundary

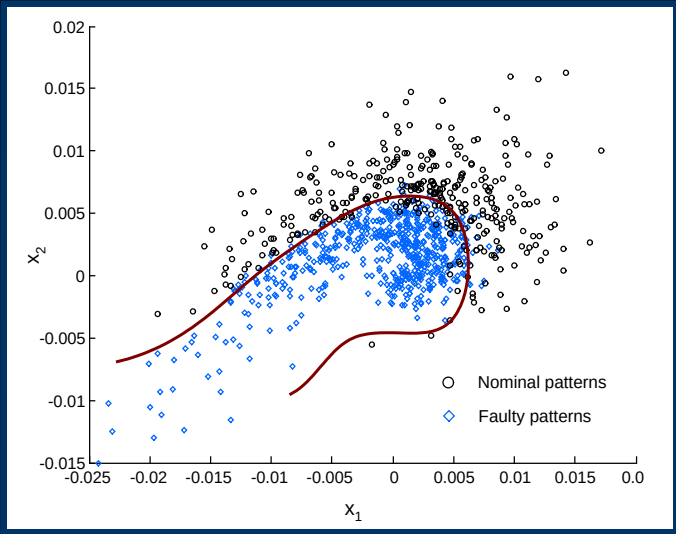

trained classifier

pass/fail label

new (untested) chip

measurement pattern

TESTING

Using a Non-Linear Neural Classifier

- Allocates a single boundary of arbitrary order
- No prior knowledge of boundary order is required
- Constructed using linear perceptrons only
- The topology is not fixed, but it grows (ontogeny) until it matches the intrinsic complexity of the separation problem

Linear Perceptron

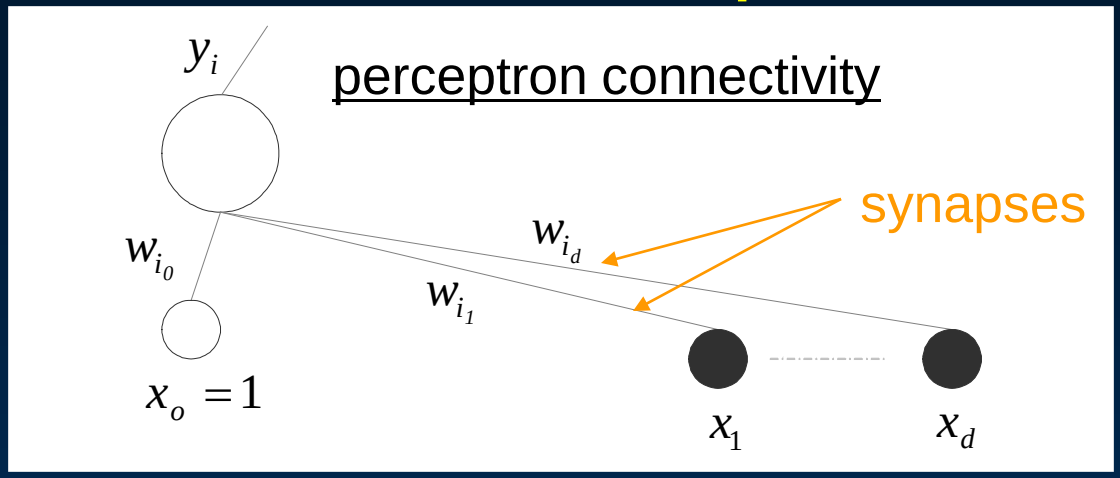

perceptron output

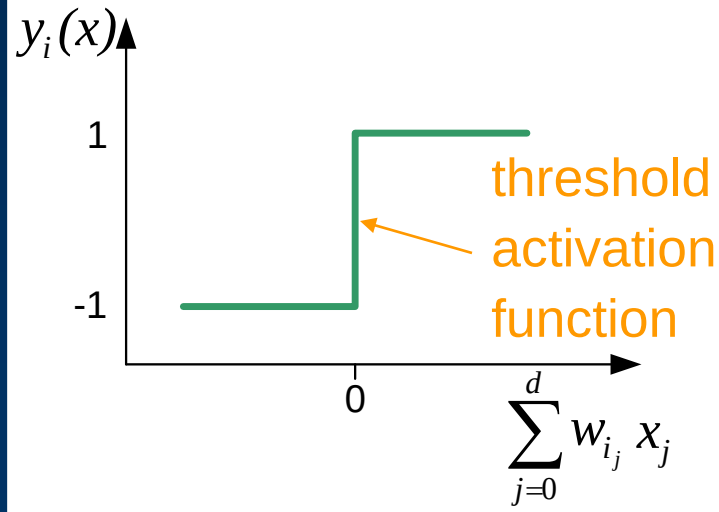

geometric interpretation

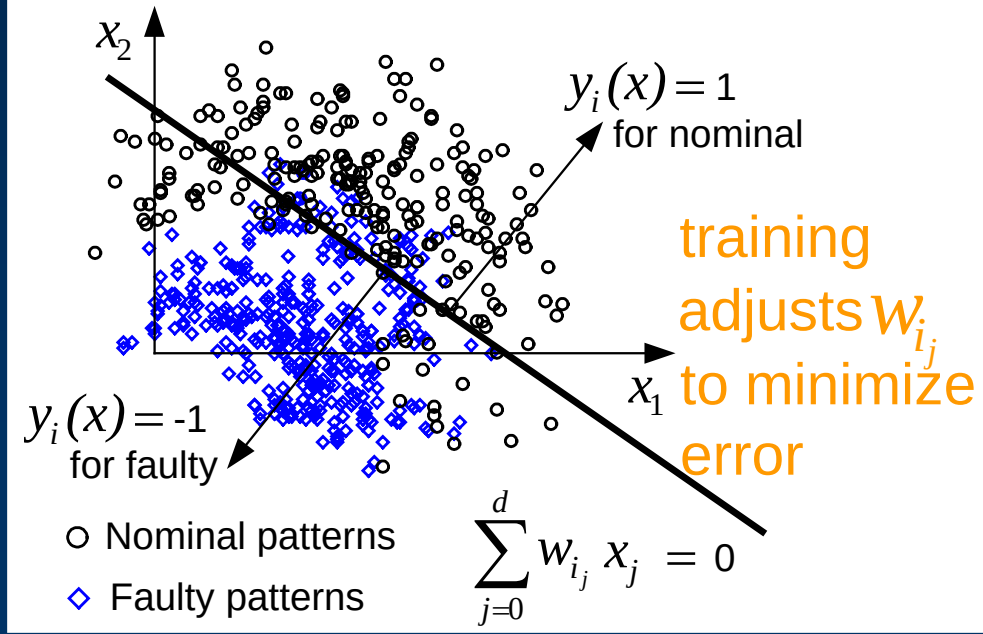

Topology of Neural Classifier

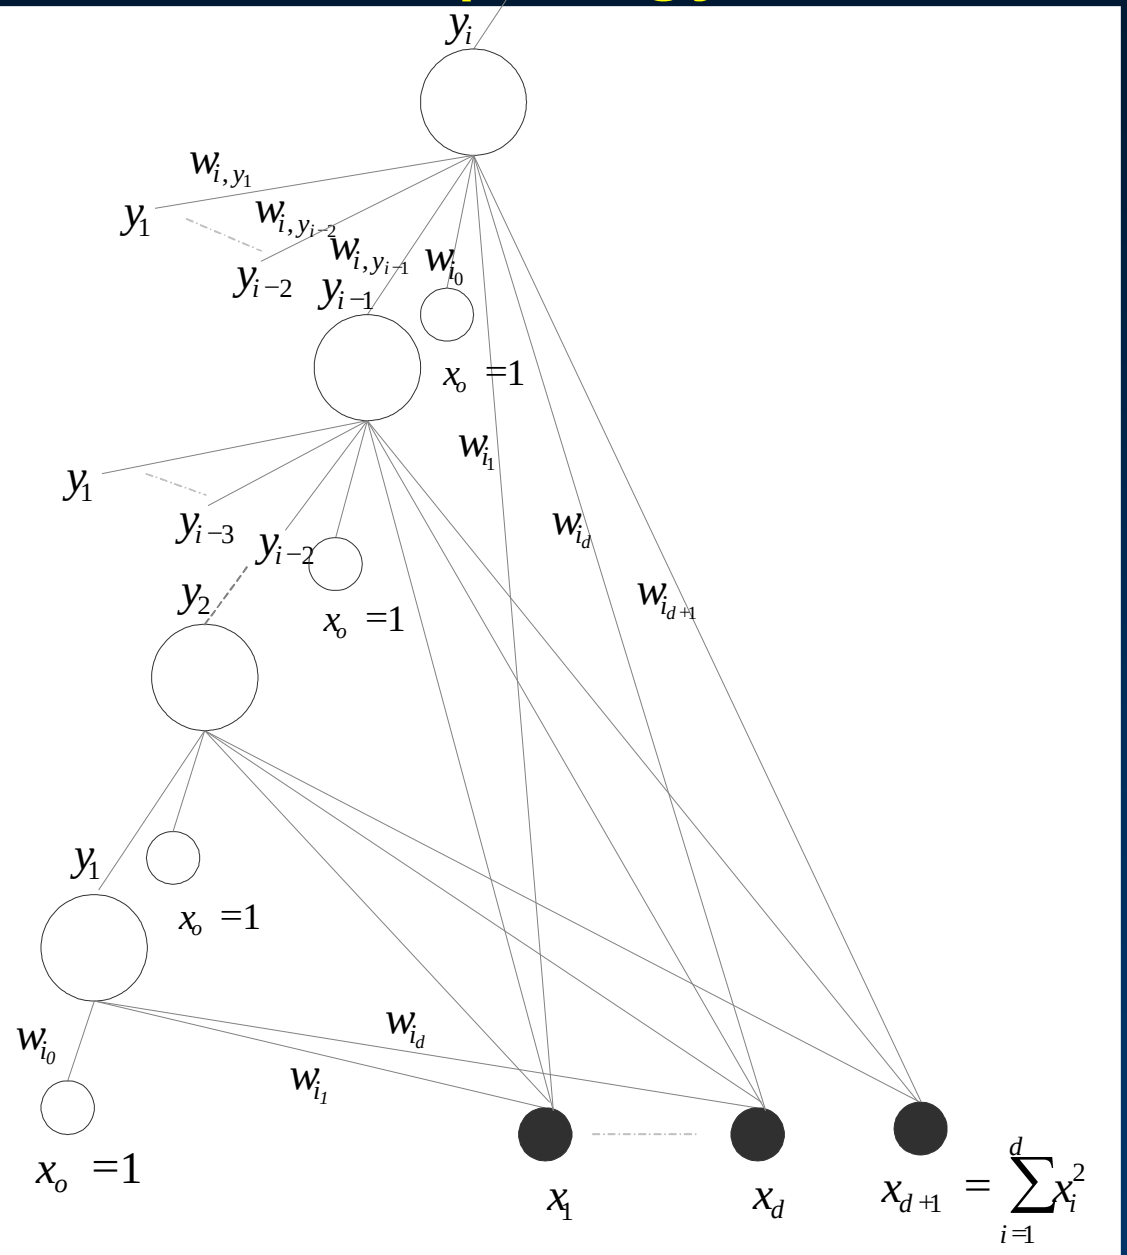

- Pyramid structure
  - First perceptron receives the pattern $\vec{x} \in R^d$
  - Successive perceptrons also receive inputs from preceding perceptrons and a parabolic pattern $x_{d+1}=\sum x_i^2$, i=1…d
- Every newly added layer assumes the role of the network output
- Non-linear boundary by training a sequence of linear perceptrons

Training and Outcome

- Weights of added layer are adjusted through the thermal perceptron training algorithm
- Weights of preceding layers do not change
- Each perceptron separates its input space linearly
- Allocated boundary is non-linear in the original space $\vec{x} \in R^d$

Theorem:

- The sequence of boundaries allocated by the neurons converges to a boundary that perfectly separates the two populations in the training set

Boundary Evolution Example (layer 0)

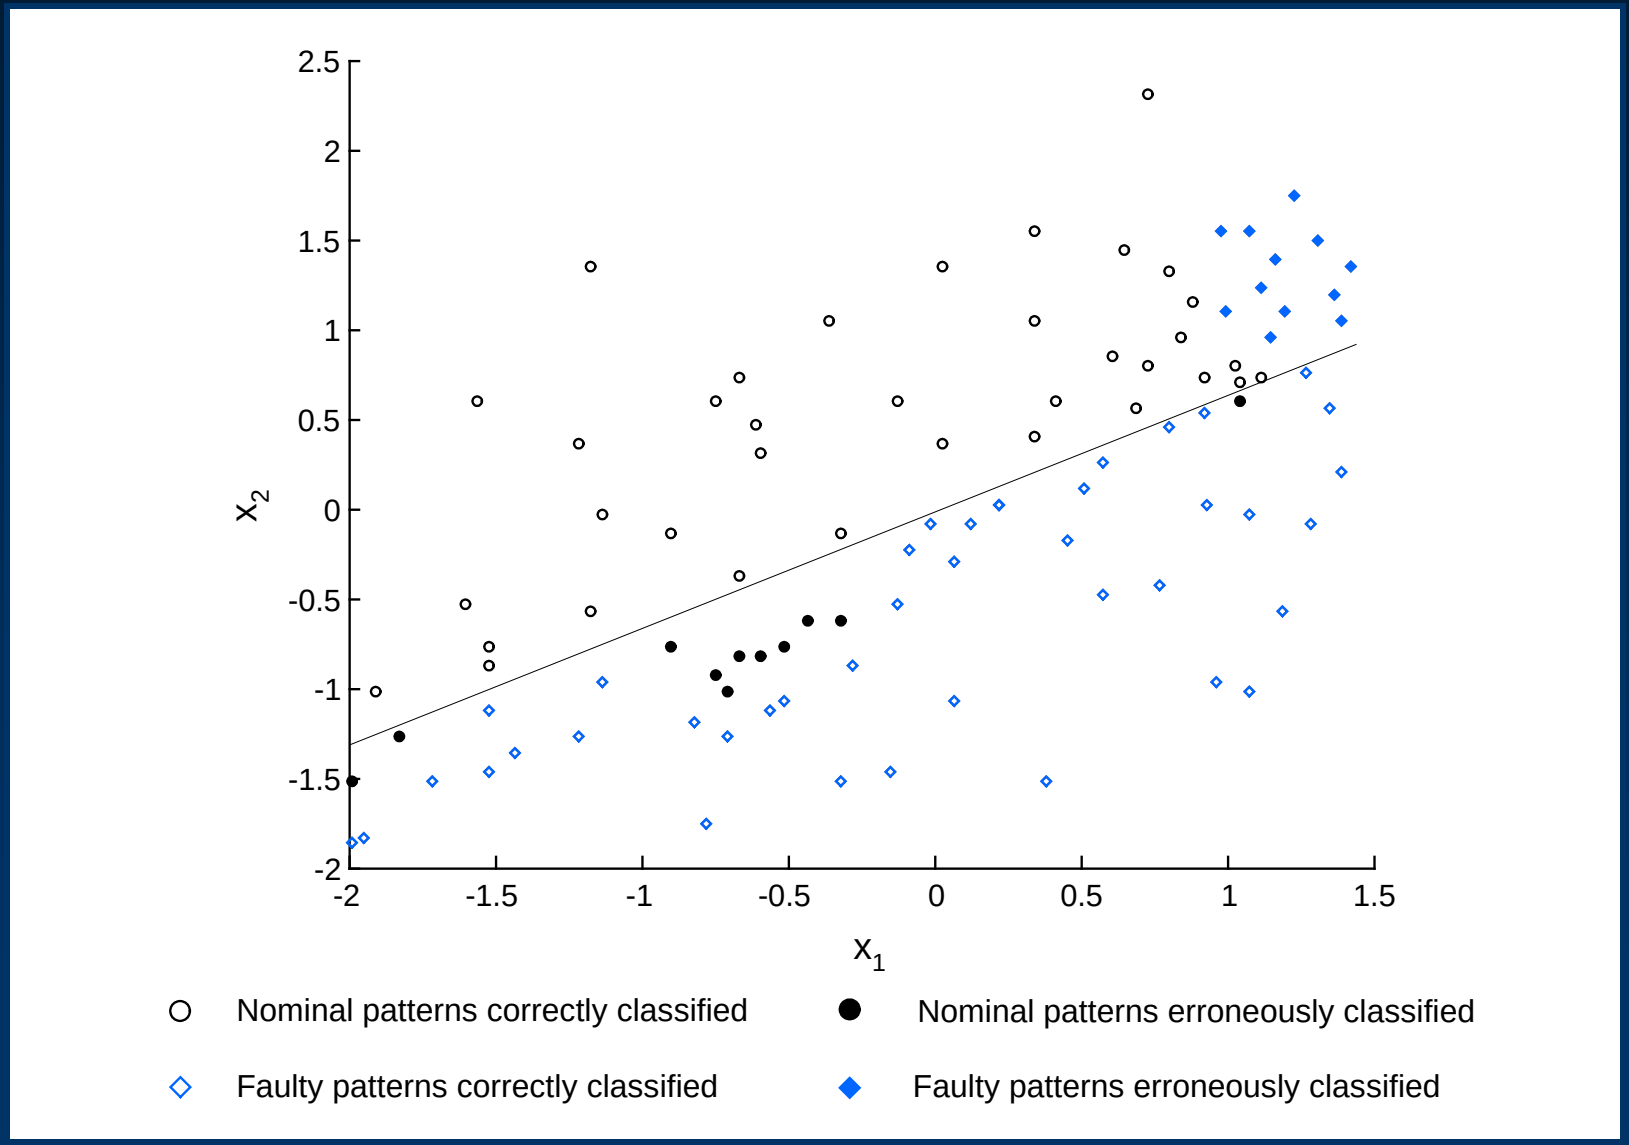

Boundary Evolution Example (layer 1)

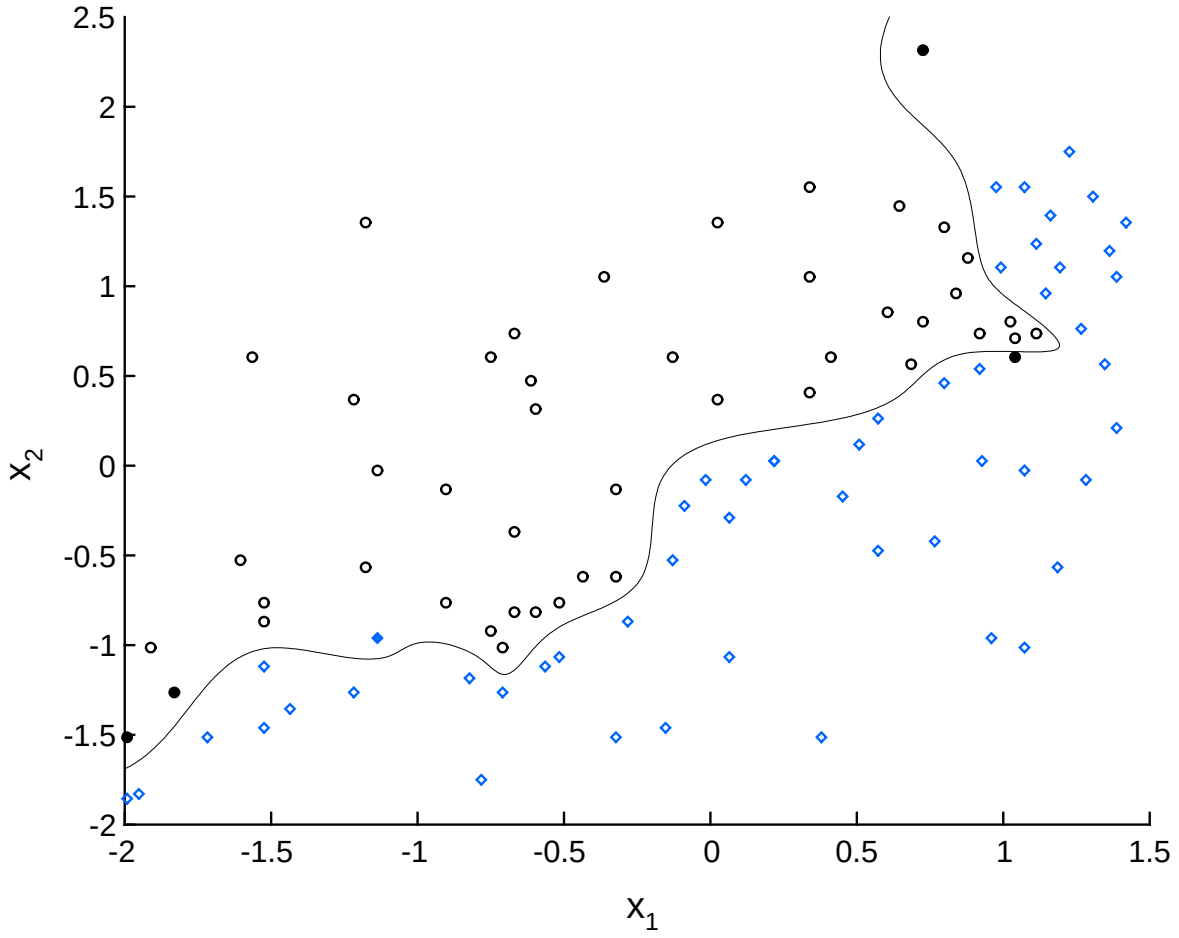

Boundary Evolution Example (layer 2)

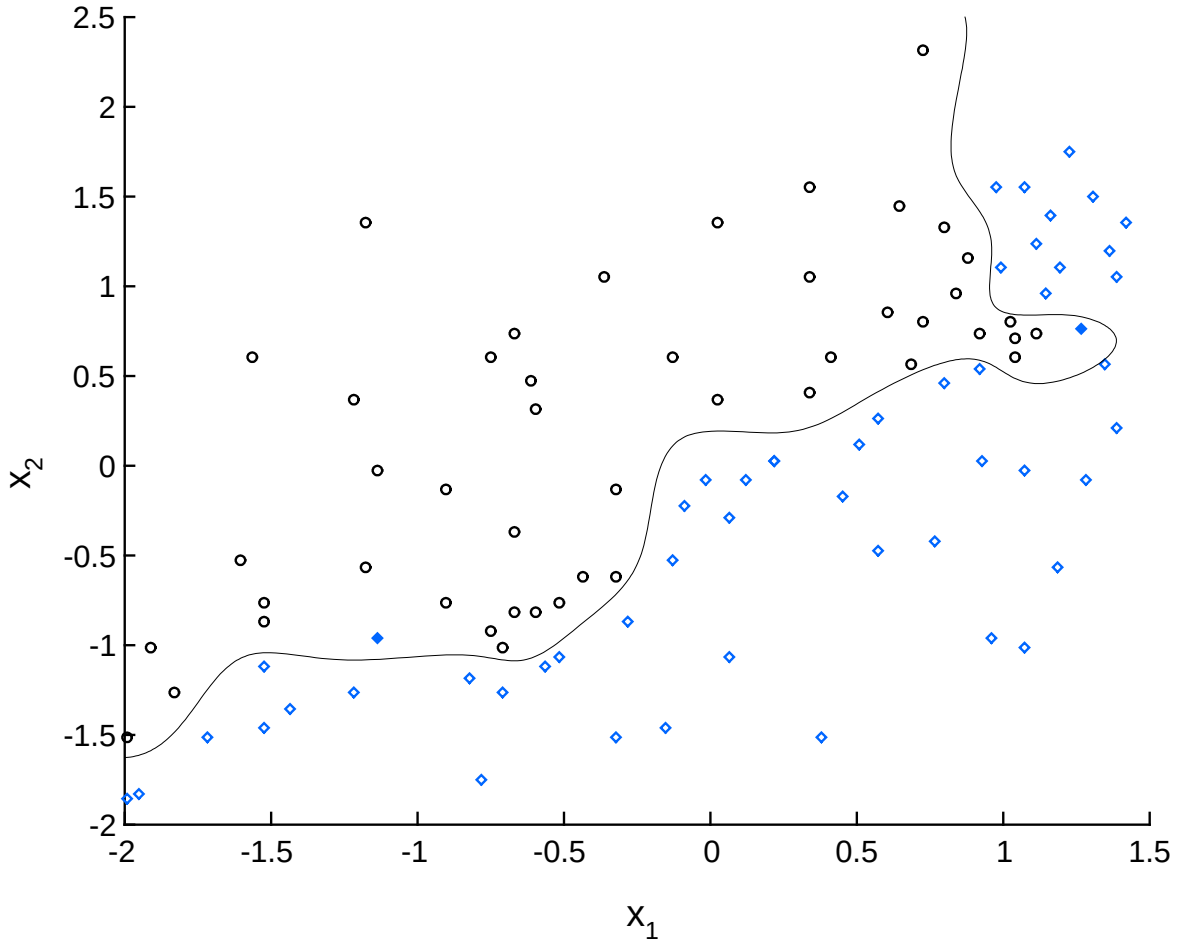

Boundary Evolution Example (output layer)

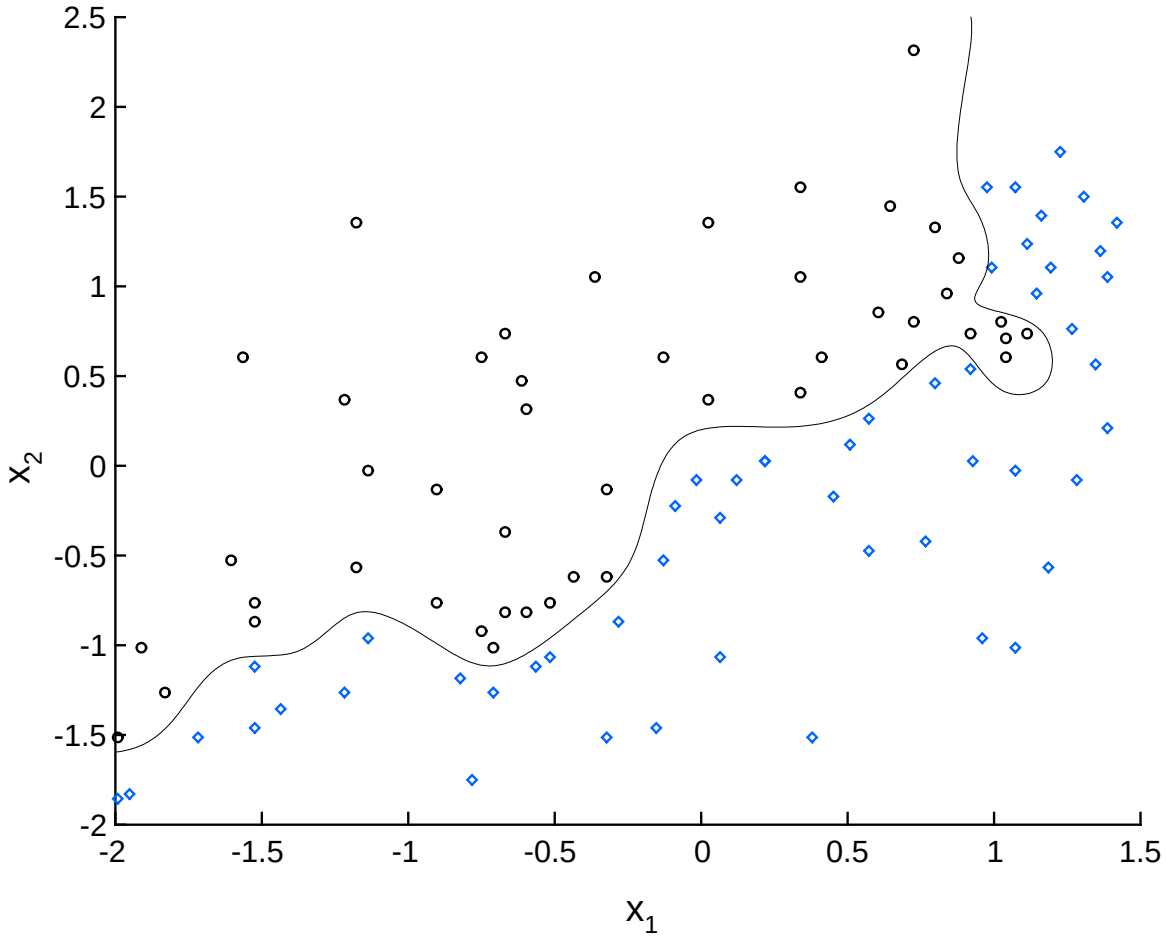

Matching the Inherent Boundary Order

Is Higher Order Always Better?

- No! The goal is to generalize
- Inflexible and over-fitting boundaries

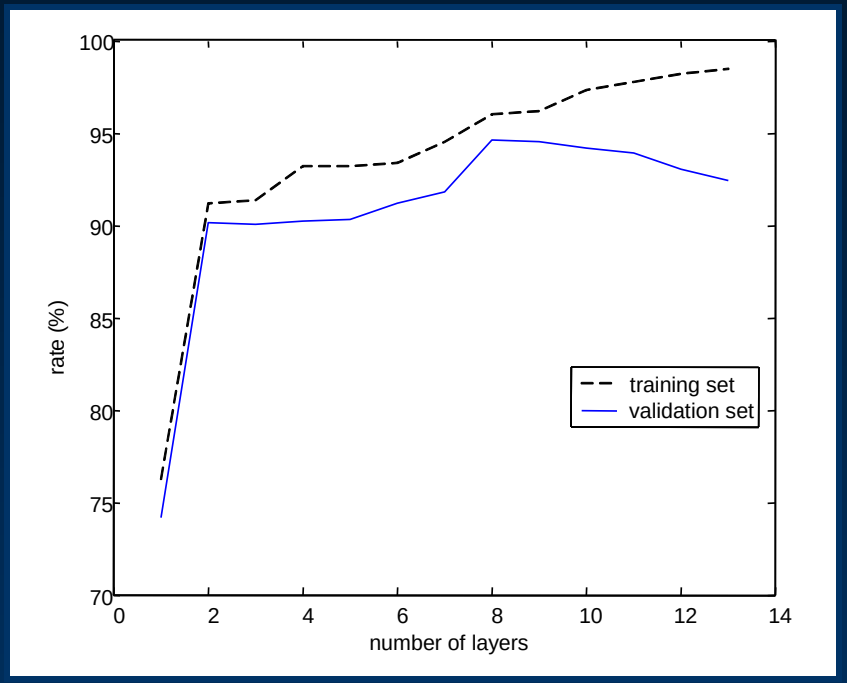

Finding The Trade-Off Point (Early Stopping)

- Monitor classification on validation set
- Prune down network to layer that achieves the best generalization on validation set
- Evaluate generalization on test set

Are All Measurements Useful?

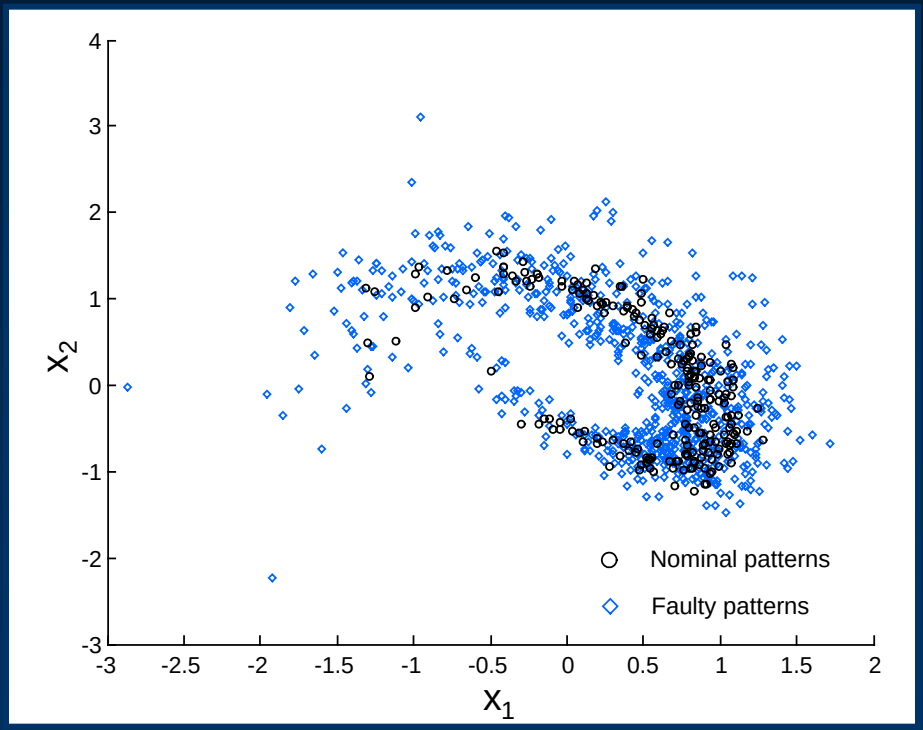

Non-Discriminatory

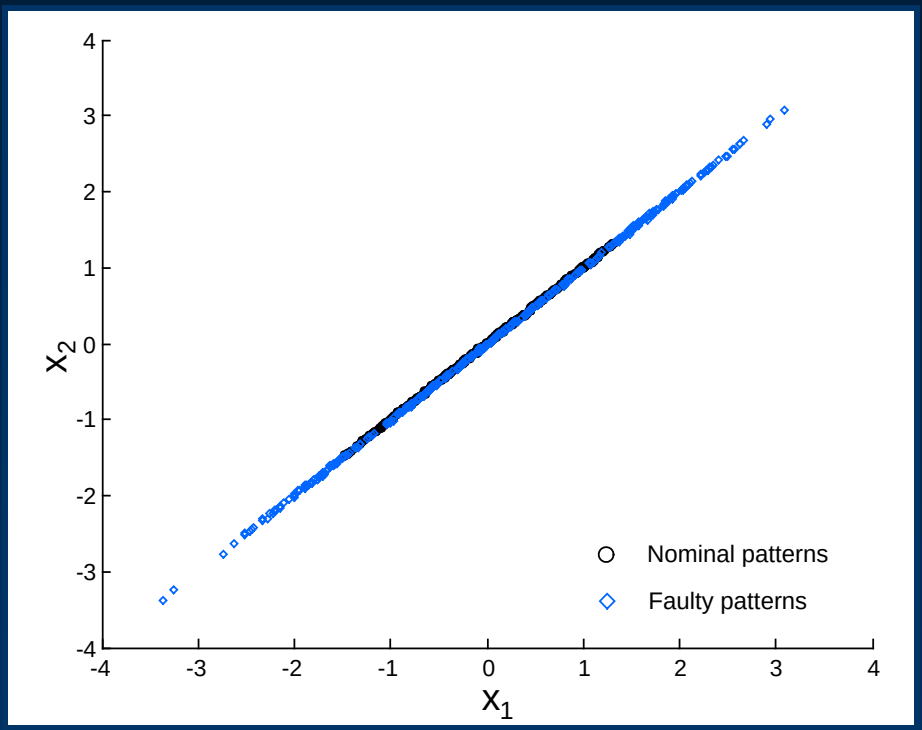

Linearly Dependent

Curse of Dimensionality

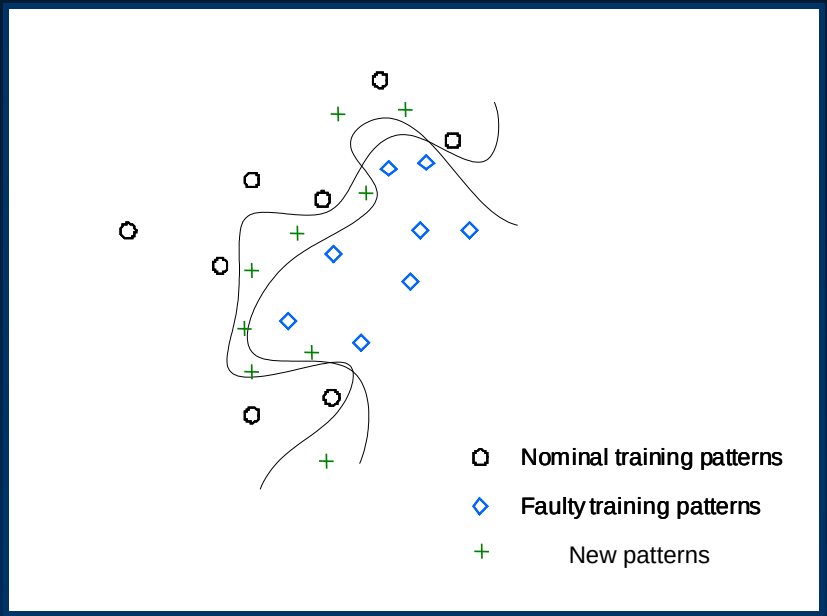

- By increasing the dimensionality we may reach a point where the distributions are very sparse
- Several possible boundaries exist – choice is random
- Random label assignment to new patterns

Genetic Measurement Selection

- Encode measurements in a bit string, with the k-th bit denoting the inclusion (1) or exclusion (0) of the k-th measurement

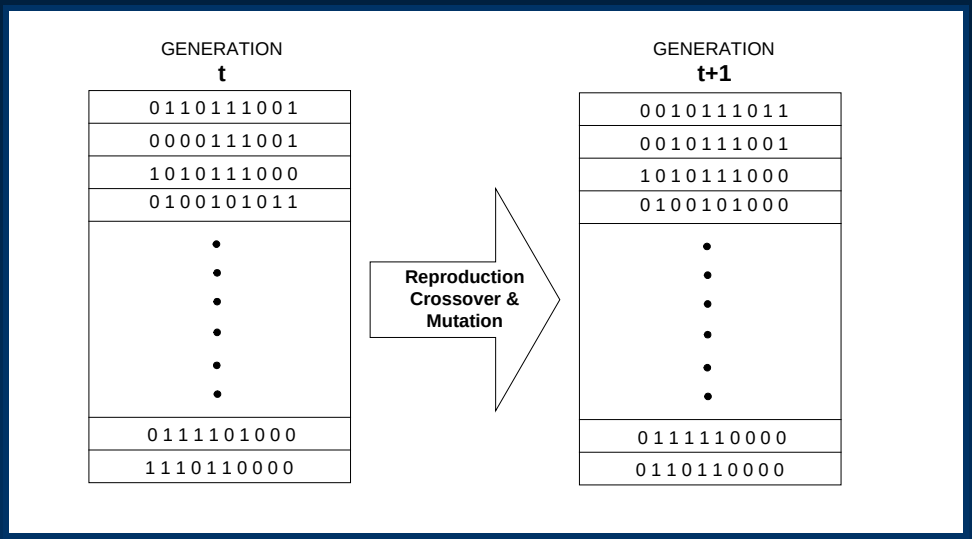

Fitness function:

NSGA II: Genetic algorithm with multi-objective fitness function reporting pareto-front for error rate (g_r) and number of re-tested circuits (n_r)

Two-Tier Test Strategy

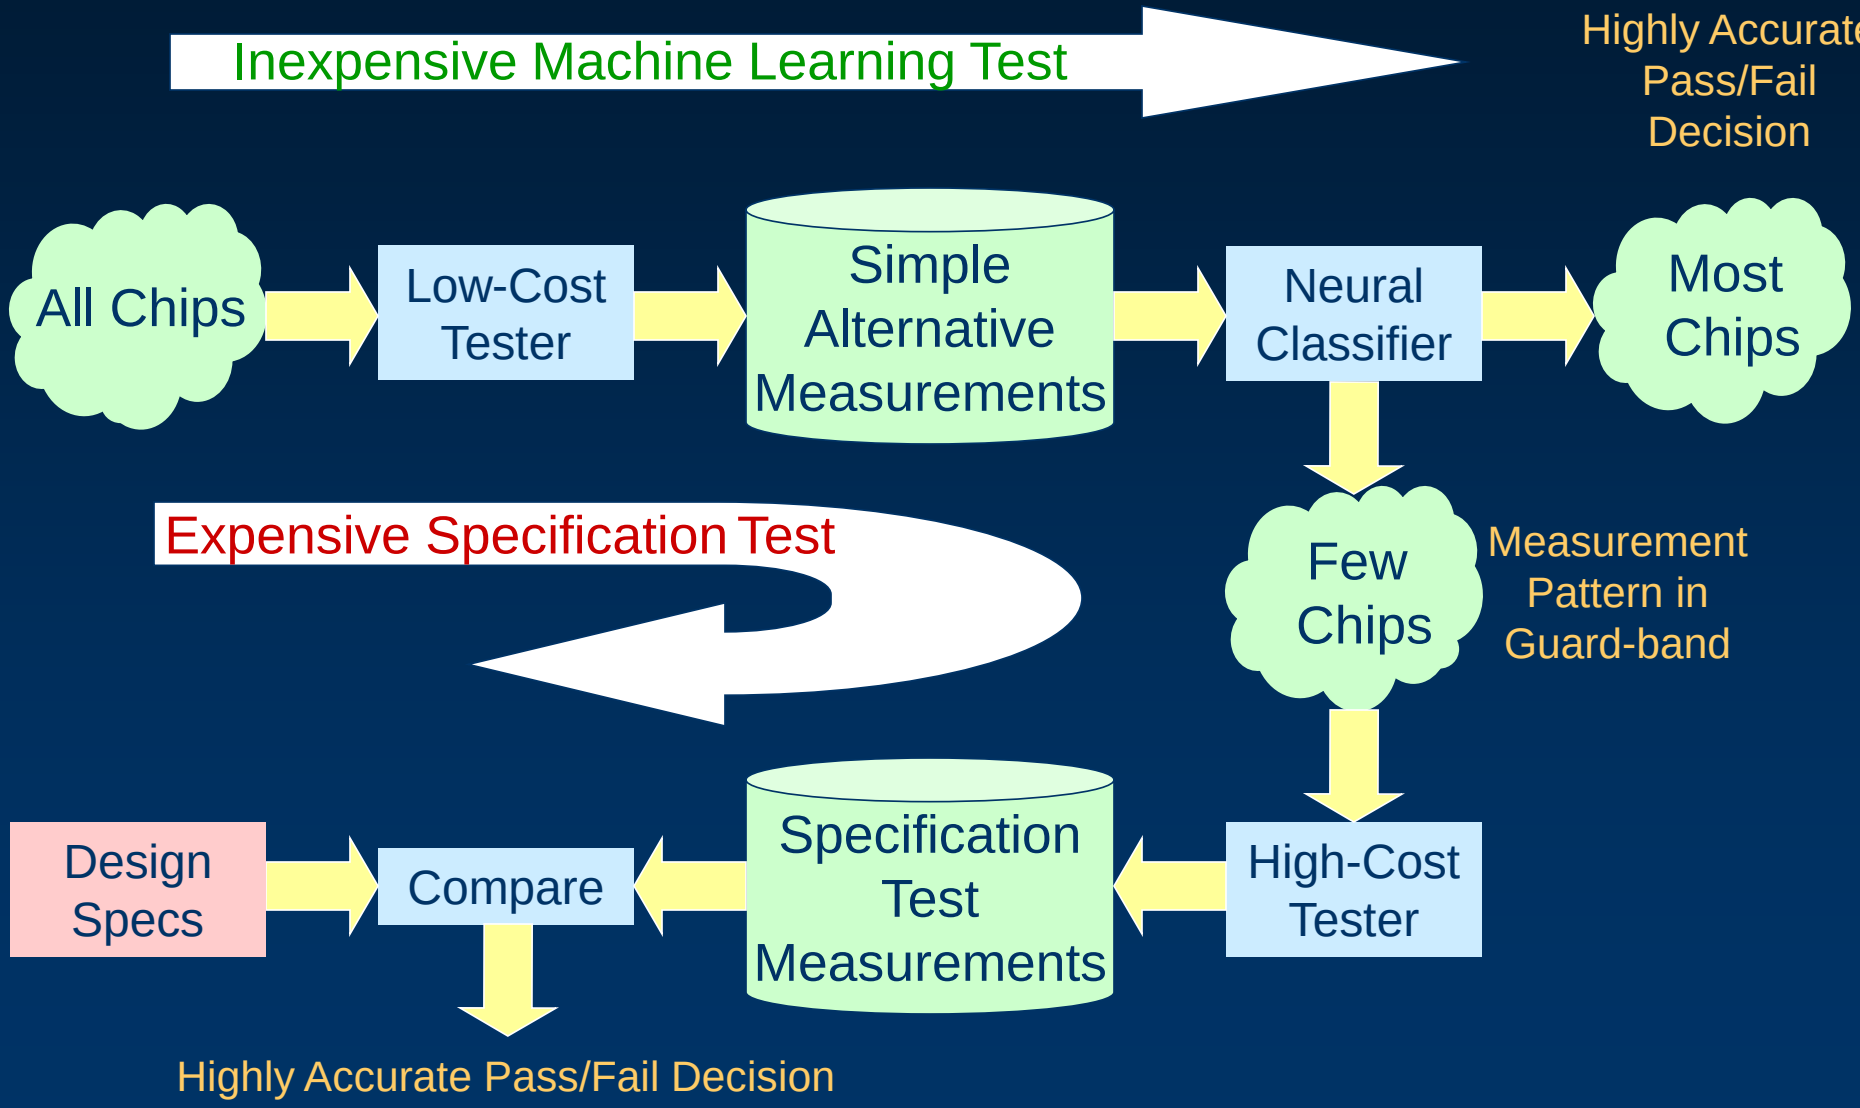

Guard-bands

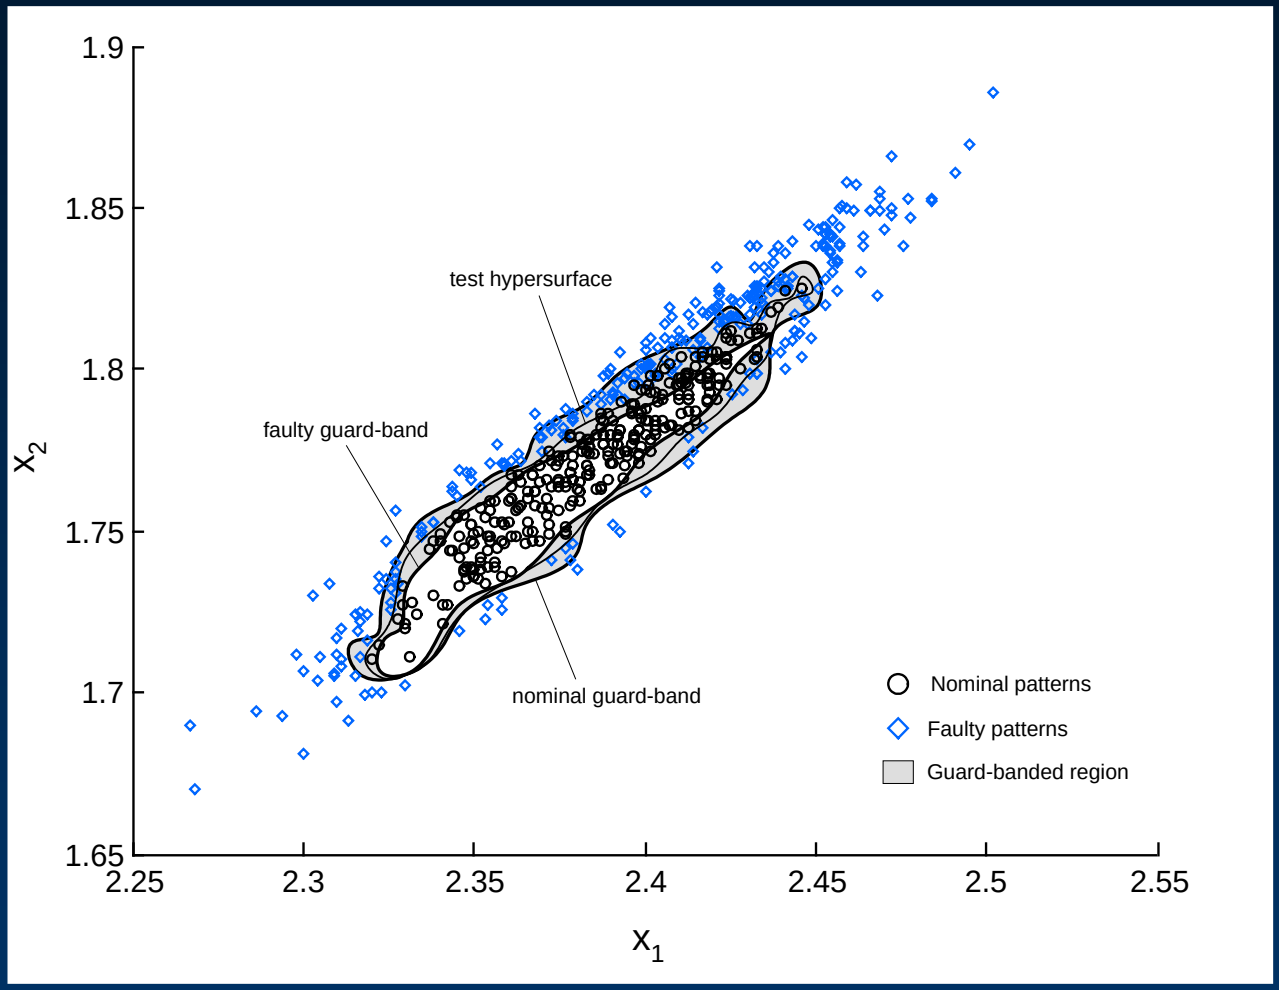

- Introduce a trichotomy in the measurement space in order to assess the confidence of the decision

Testing Using Guard-bands

Re-test via specification test

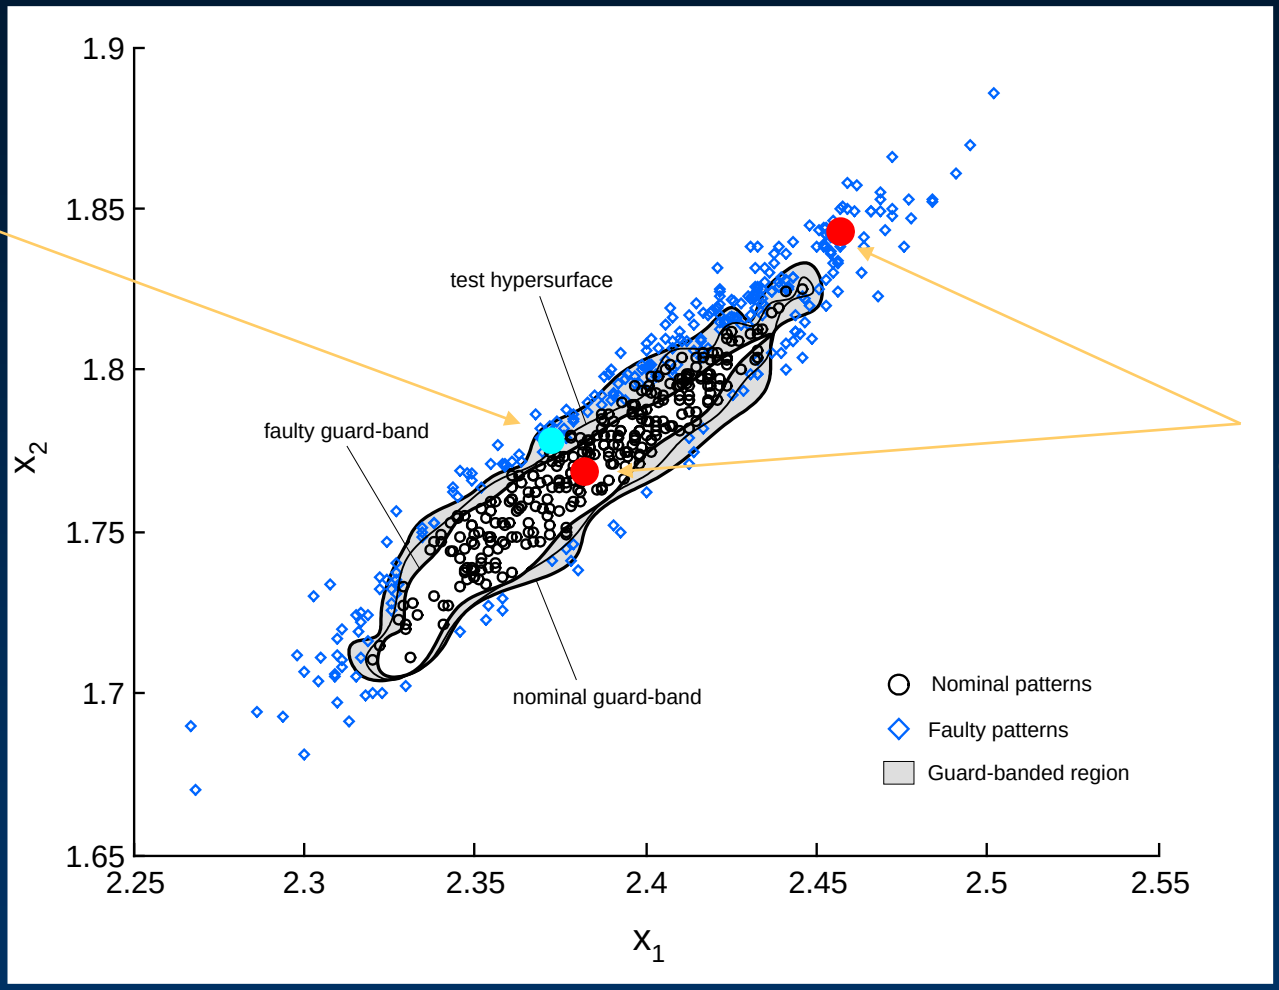

Classify chip to dominant class

- Examine measurement pattern with respect to the guard-bands
- Identify circuits that need to be re-tested via specification tests

Guard-band Allocation

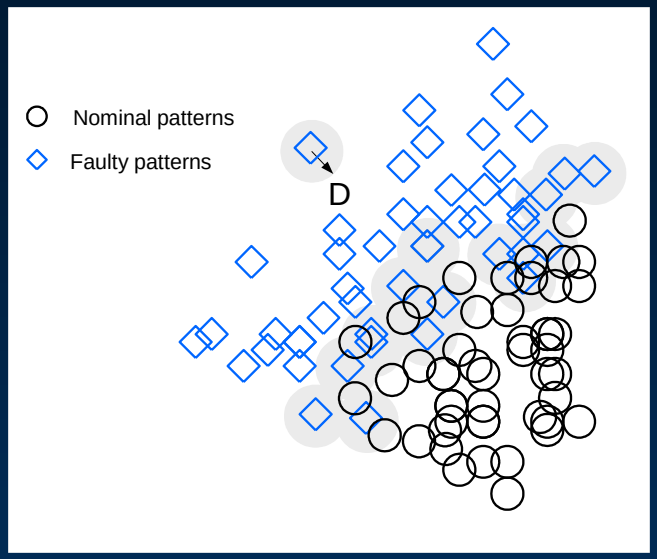

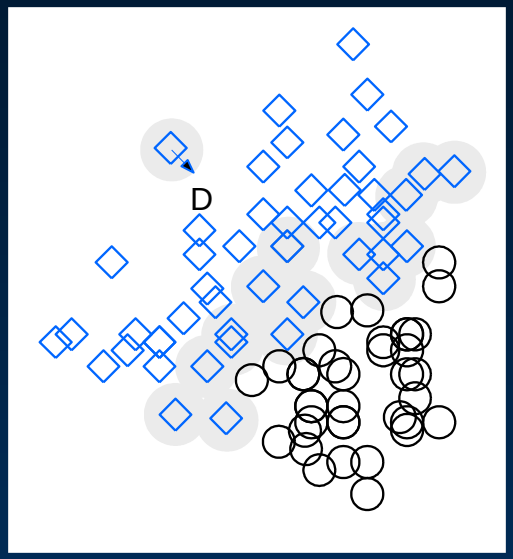

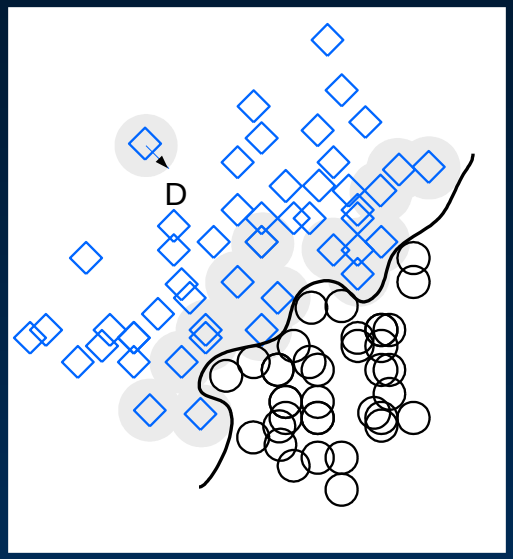

- Guard-bands are allocated separately

- Identify the overlapping regions $$D = \frac{1}{N_f} \sum_{i \in C_f} \min_{j \in C_n} \| \vec{x}^i - \vec{x}^j \|$$

- Clear the overlapping regions $$\| \vec{x}^i - \vec{x}^j \| < D$$

- The ontogenic classifier allocates the guard-band

Guard-band Allocation

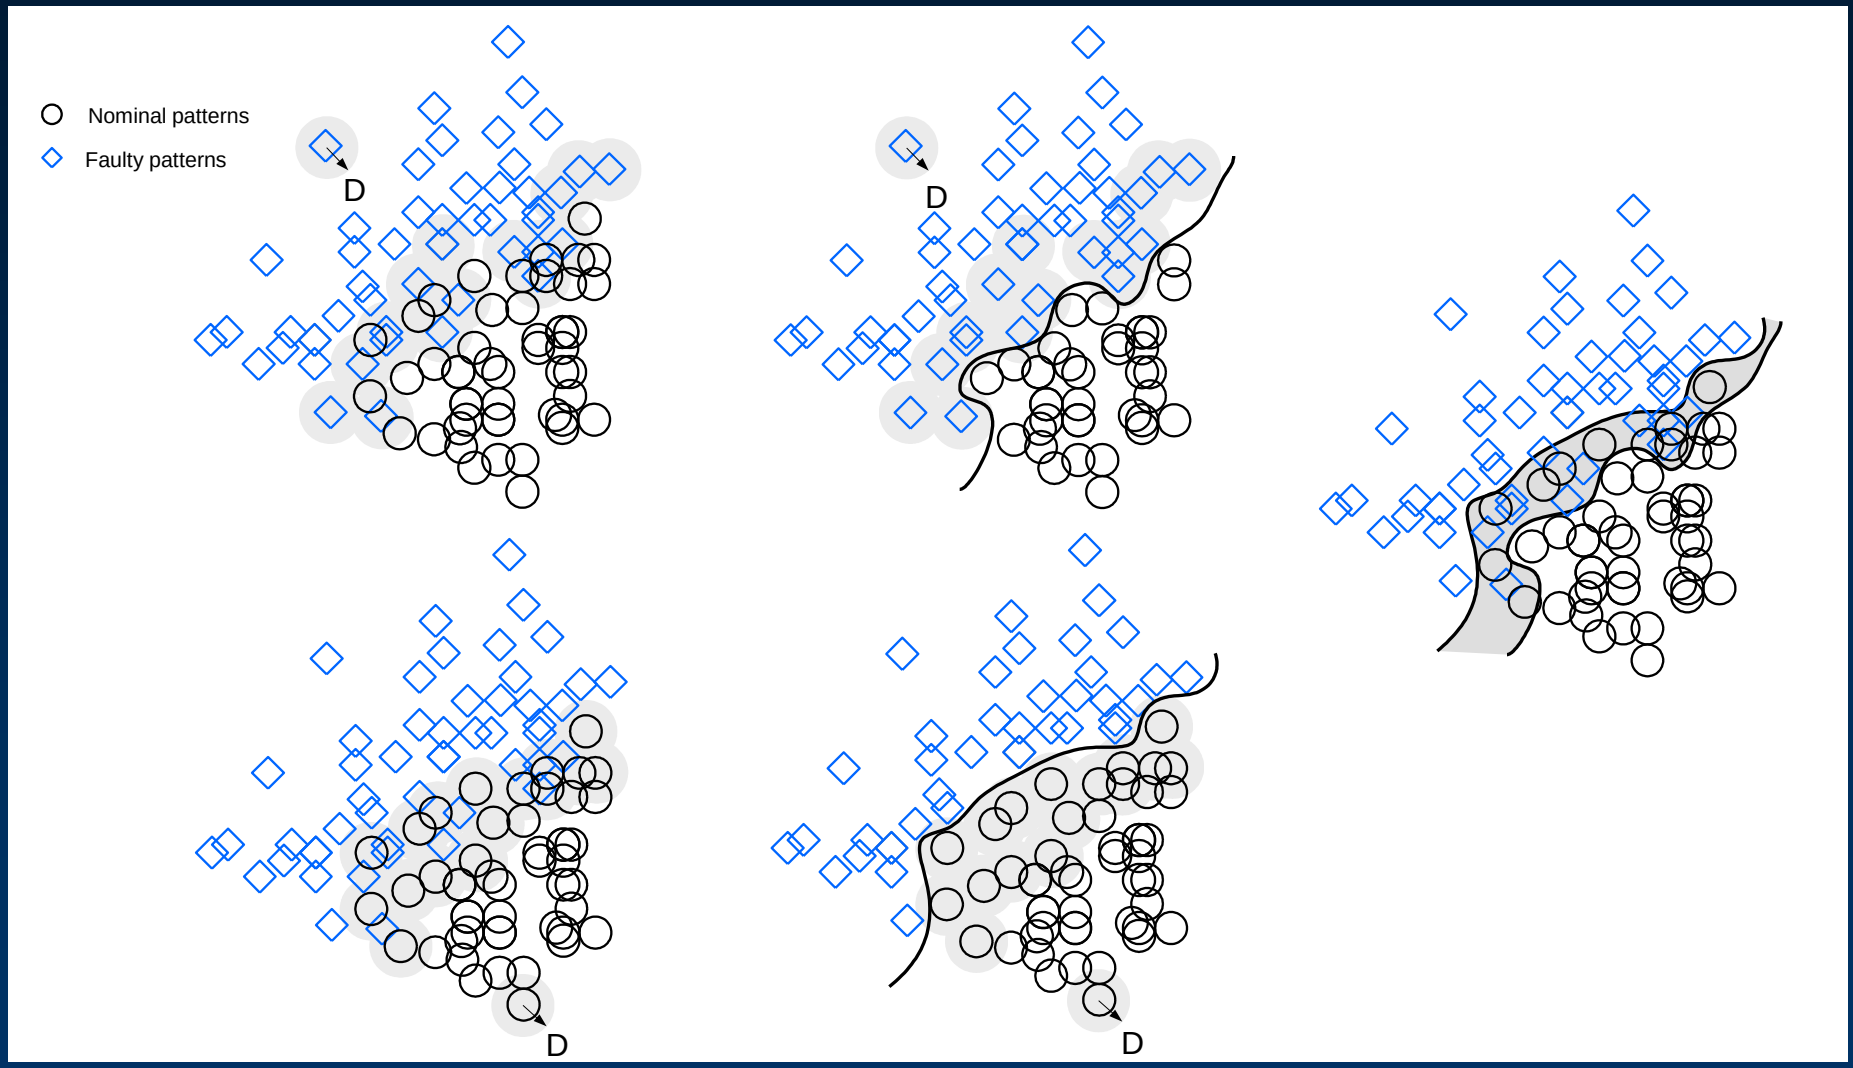

- The guard-banded region can been varied by controlling D

Experiment 1: Switch-Capacitor Filter

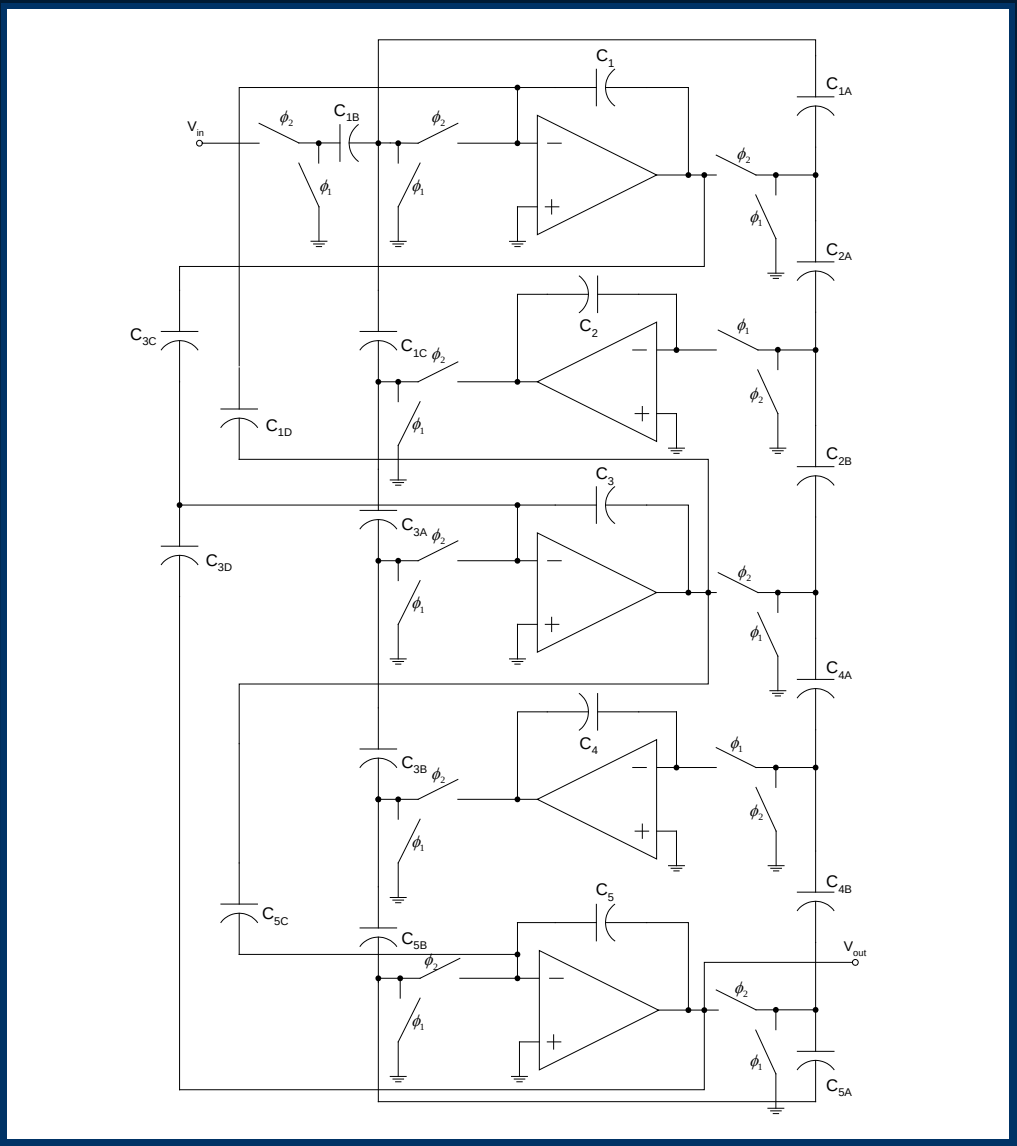

Specifications considered

- Ripples in stop- and pass-bands
- Gain errors
- Group delay
- Phase response
- THD

Experiment Setup

- N=2000 instances
- N/2 assigned to the training set, N/4 to the validation set and N/4 to the test set

White-Noise Test Stimulus

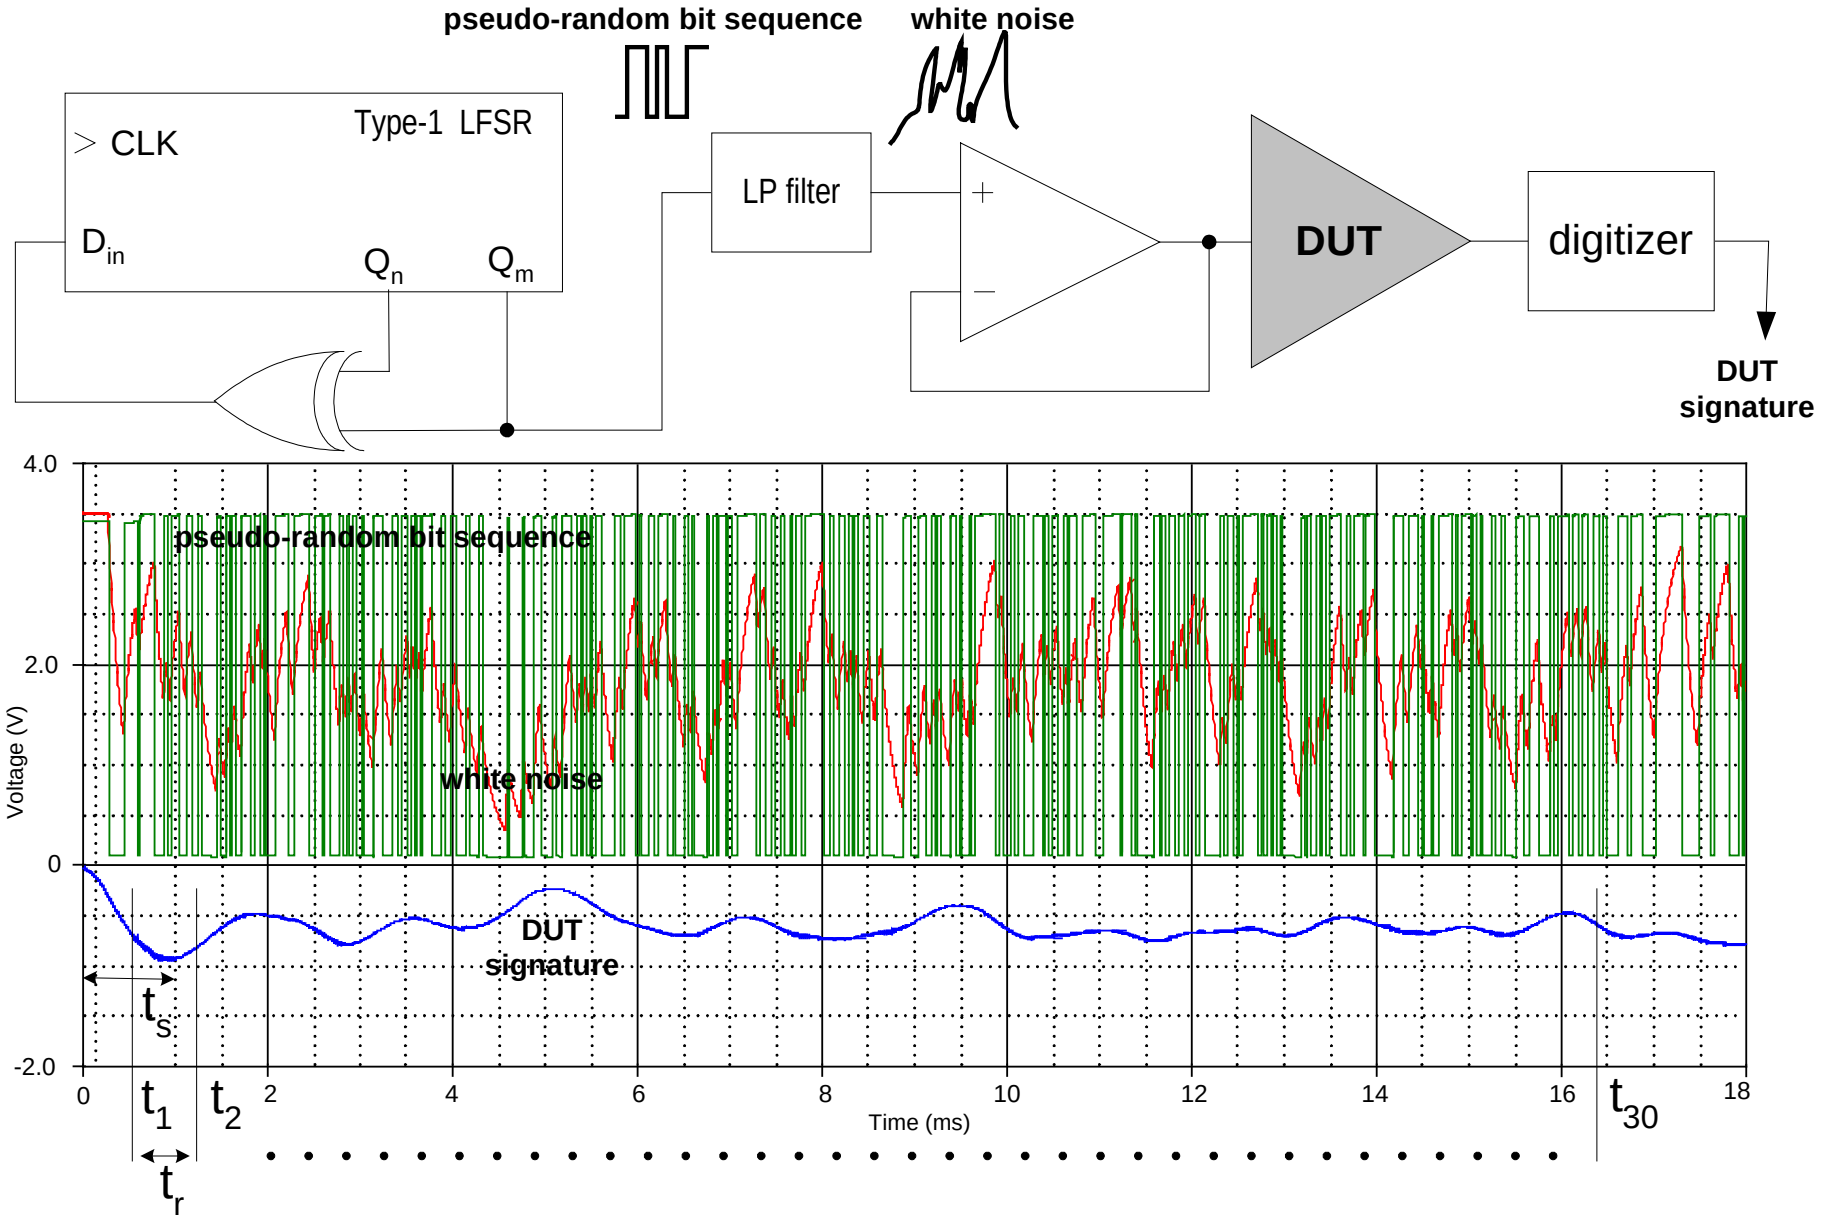

Test Time vs. Test Accuracy Trade-Off

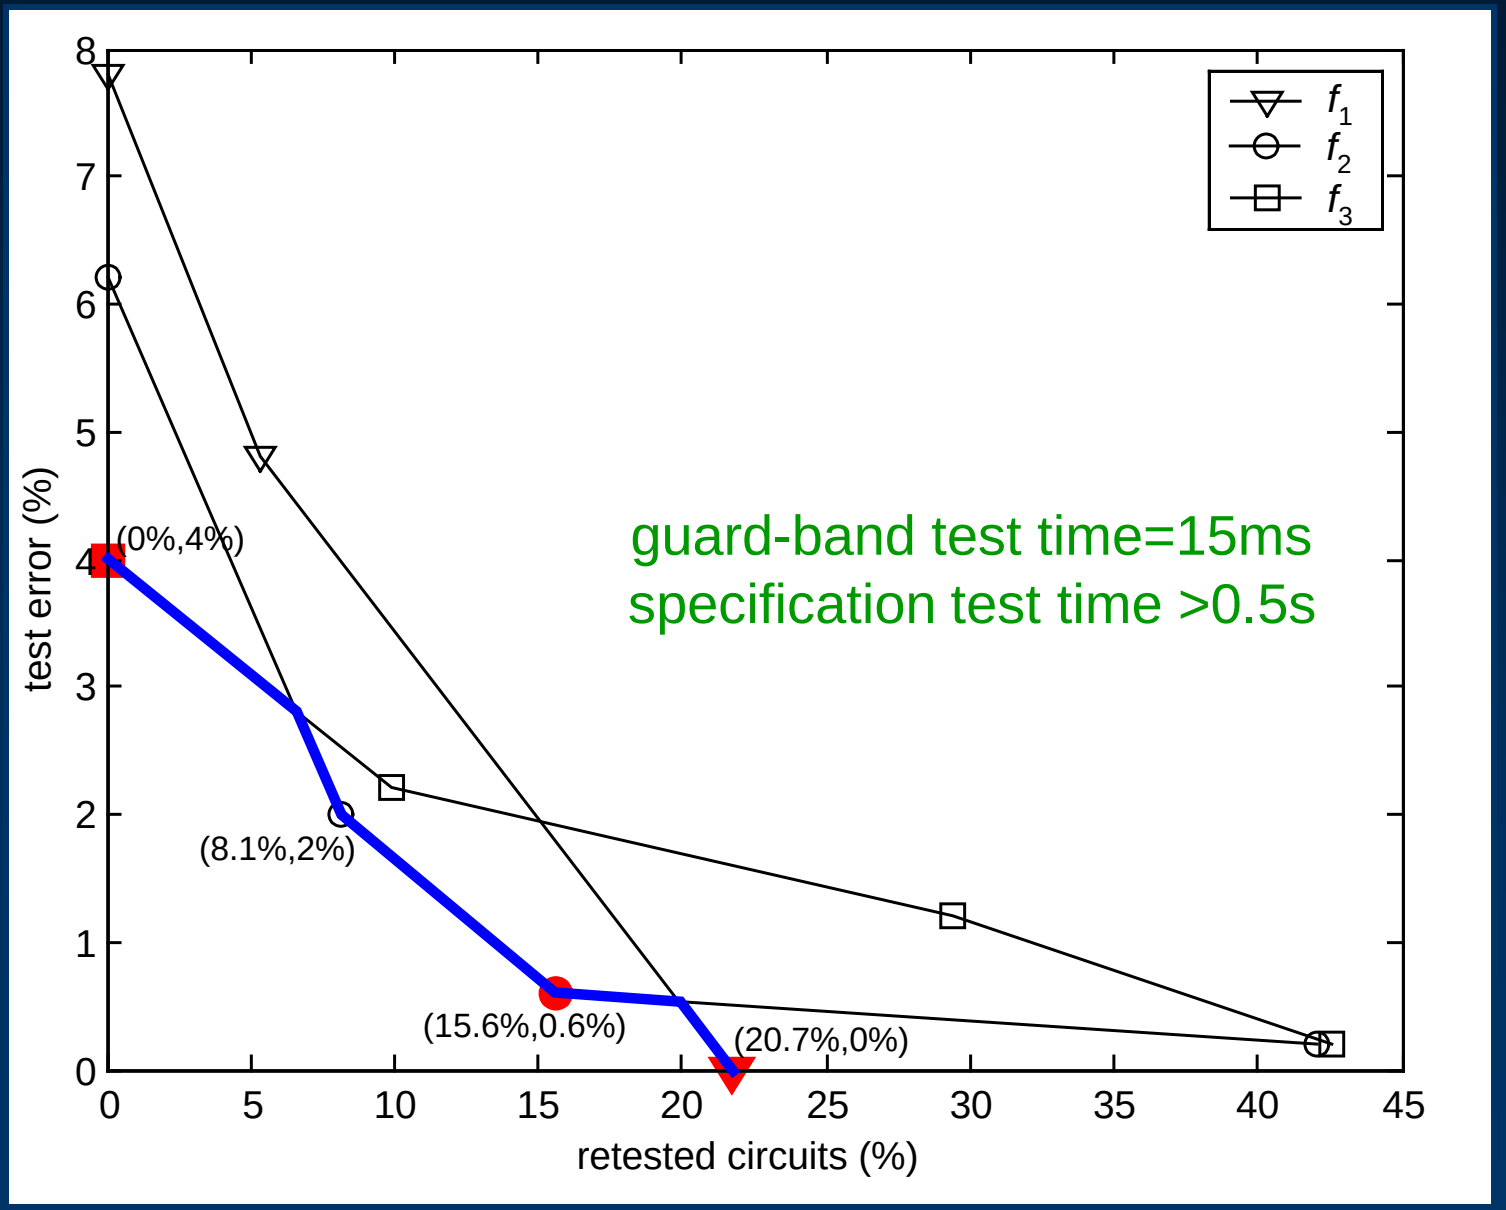

Experiment 2: RFMD2411

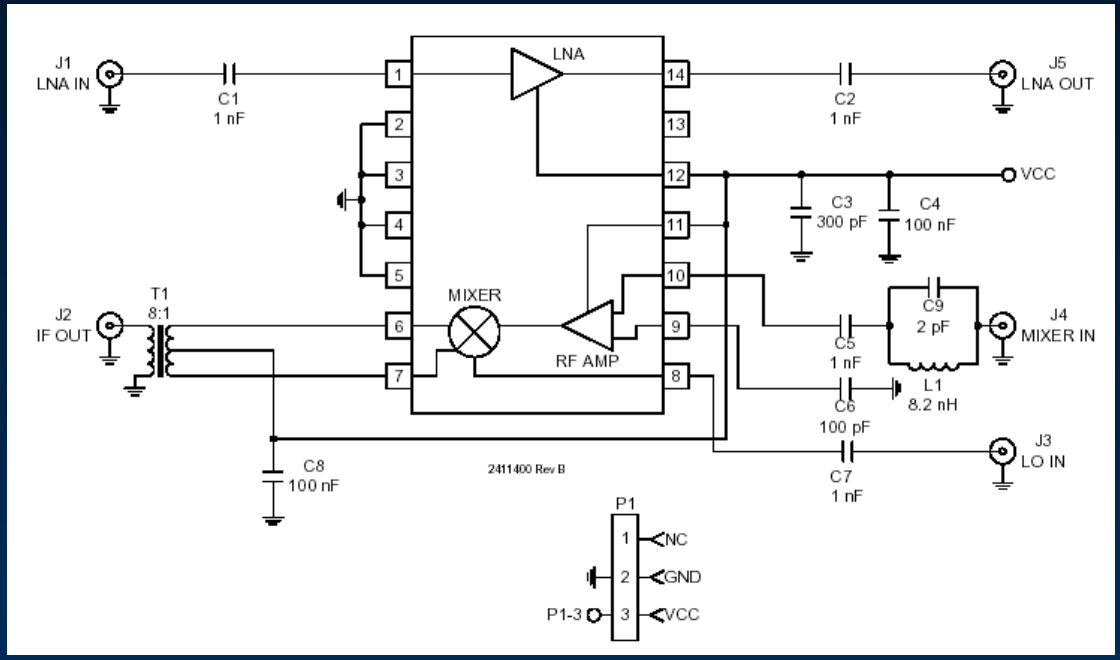

- 541 devices (training set 341, validation set 100, test set 100)
- Considered 13 specs @850MHz (7 different test configurations)

	Gain	3rd Order Intercept	Noise Figure	Other S Parameters		
LNA	X	X	X	X	X	X
Mixer	X	X	X	X	-	-
LNA + Mixer (Cascade)	X	X	X	-	-	-

Test Configuration

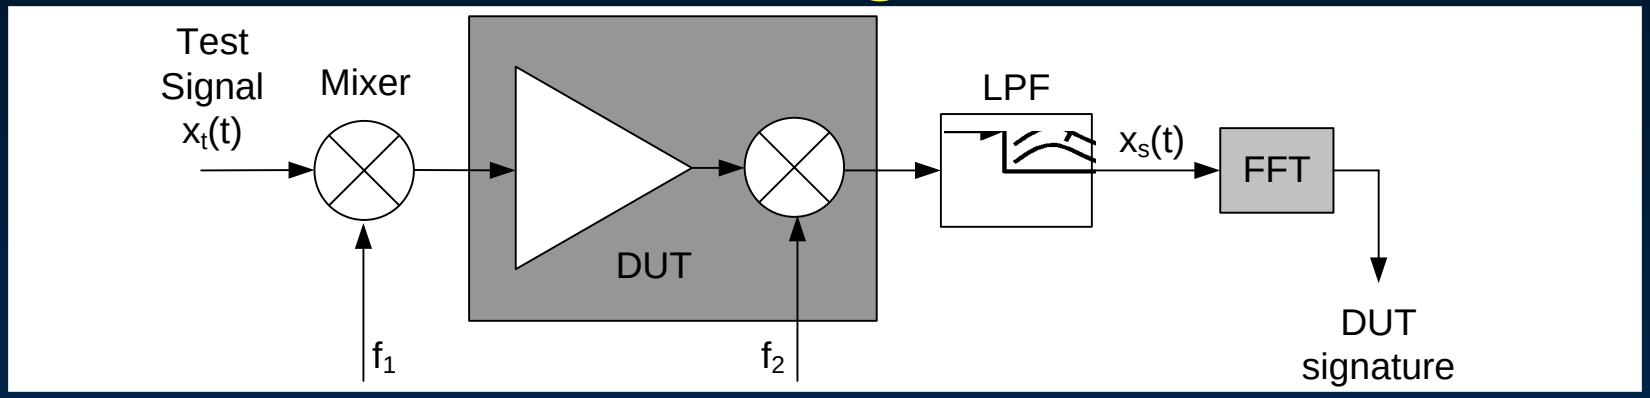

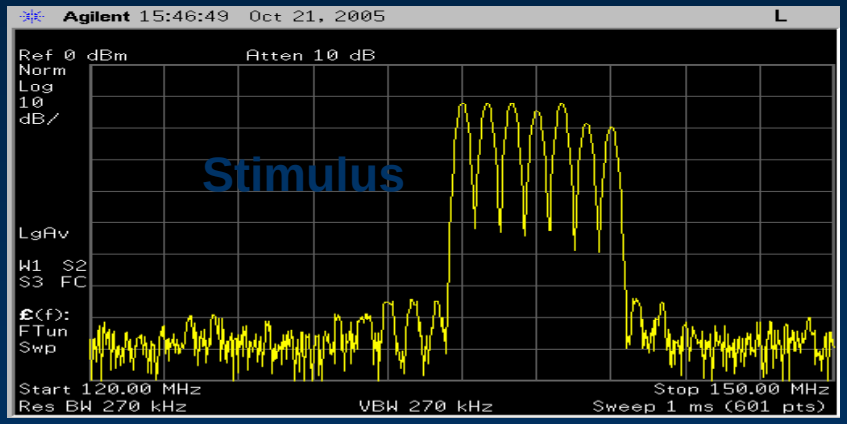

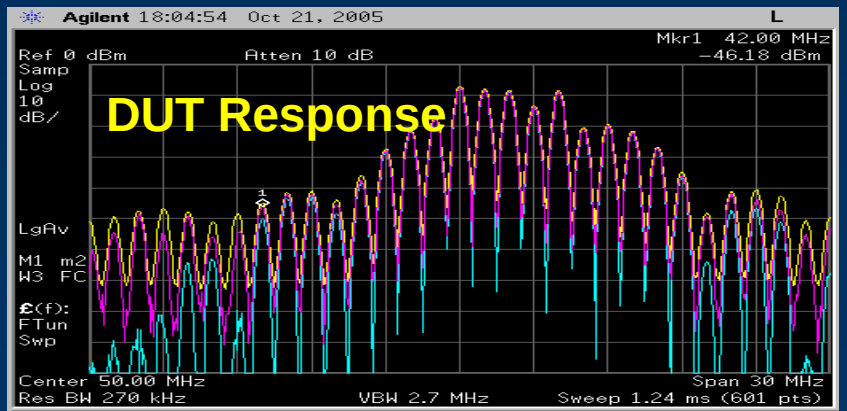

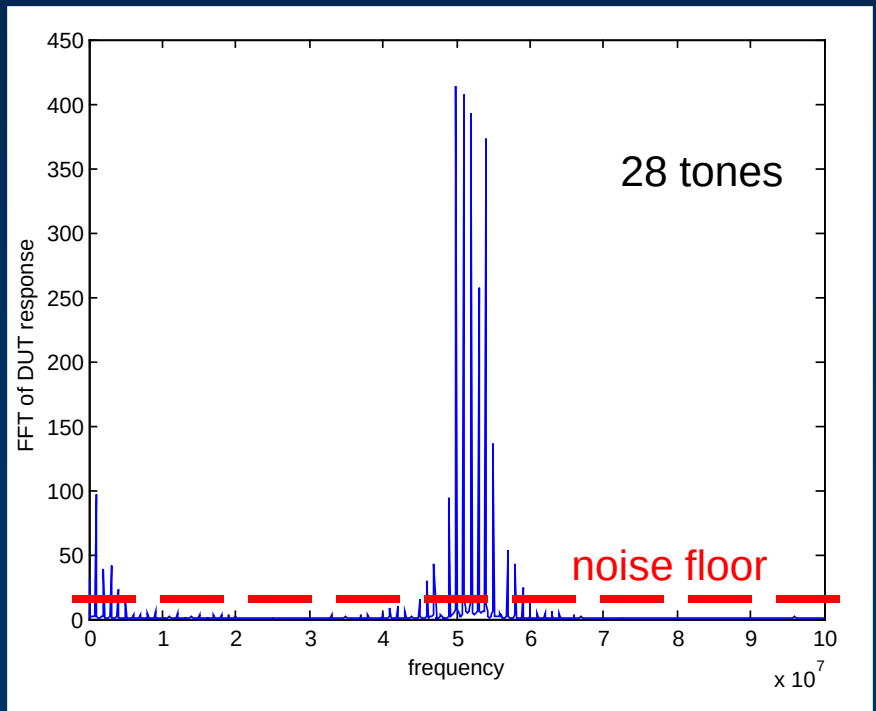

Test Time vs. Test Accuracy Trade-Off

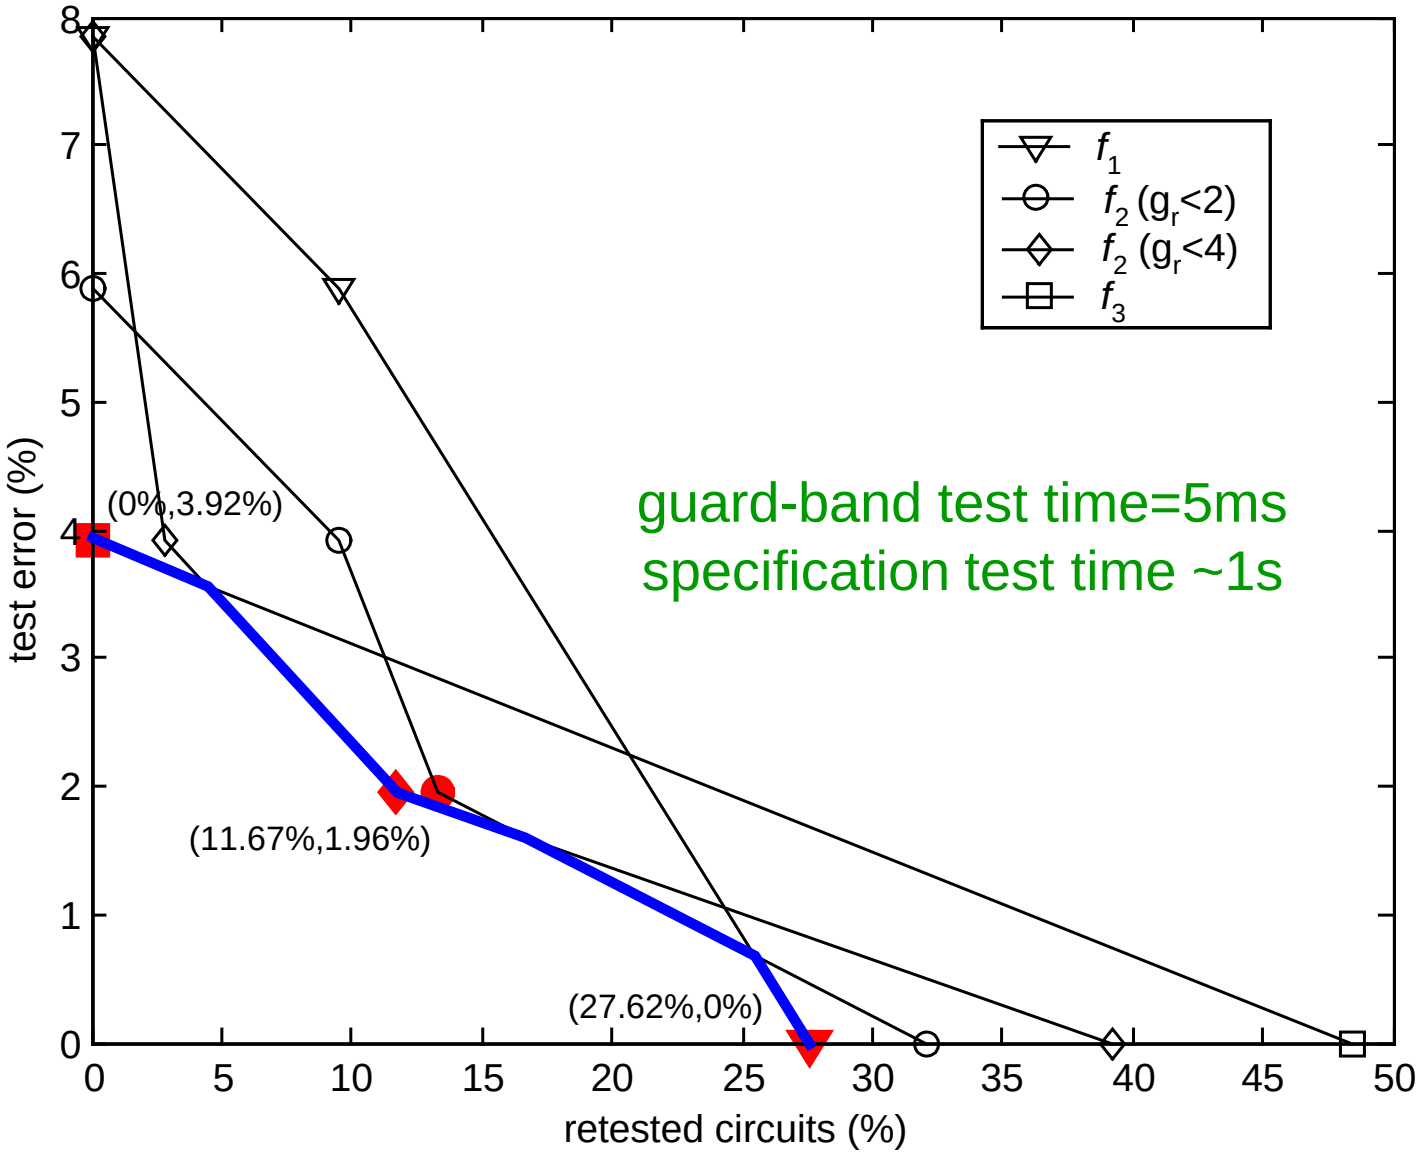

Analog/RF Specification Test Compaction

Non-disruptive application:

- Instead of “other, alternate measurements” use existing inexpensive specification tests to predict pass/fail for the chip and eliminate expensive ones

Case-Study:

- RFCMOS zero-IF down-converter for cell-phones
- Fabricated at IBM, in production until late 2008
- Data from 944 devices (870 pass – 74 fail)
- 136 performances (65 non-RF – 71 RF)
- Non-RF measurements assumed much cheaper than RF, aim at minimizing number of measurements

Setup and Questions Asked

Setup:

- Split 944 devices into training and test set (472 devices each, chosen uniformly at random.
- Repeat 200 times and average to obtain statistically significant results

Question #1:

- How accurately can we predict pass/fail of a device using only a subset of non-RF measurements?

Question #2:

- How does this accuracy improve by selectively adding a few RF measurements to this subset?

Results

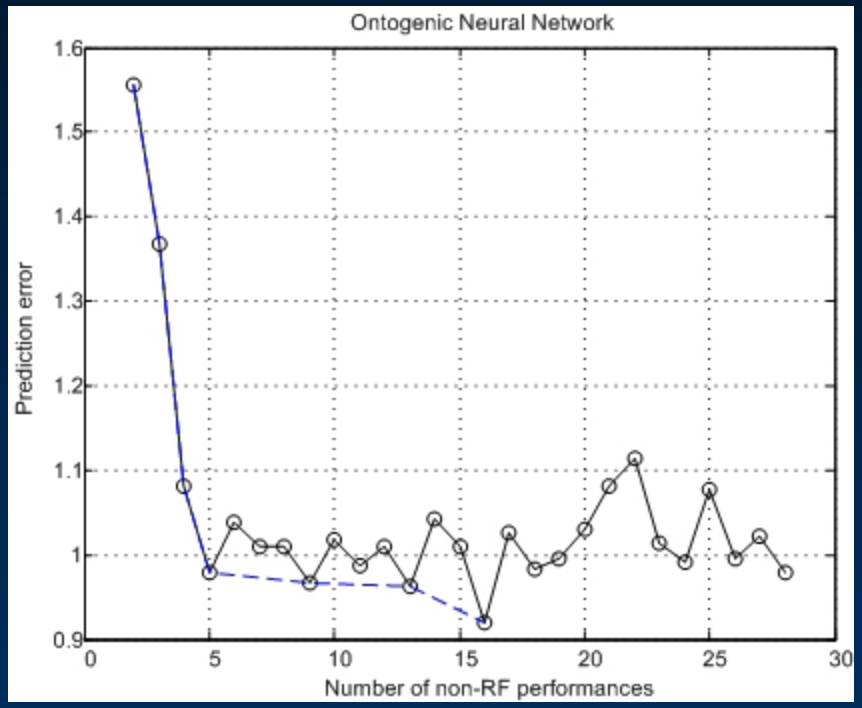

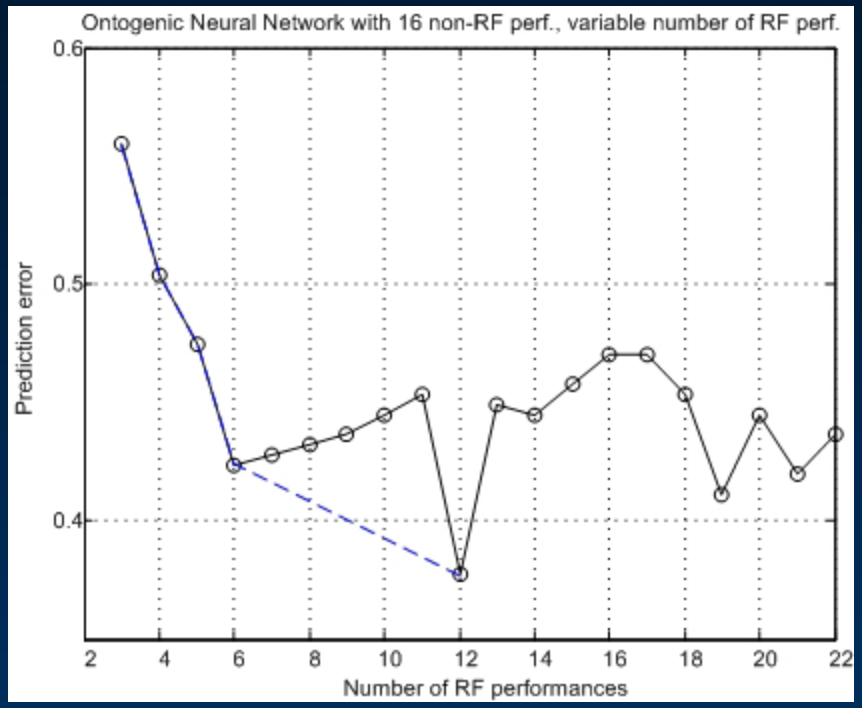

Only non-RF performances:

- 16 measurement suffice to predict correctly over 99% of the devices (in our case 4 out of 472 are on average mispredicted

Adding RF performances:

- By adding to these 16 non-RF performances 12 RF performances, error drops to 0.38% (i.e. 1 out of 472 devices mispredicted)

Continuation of Case-Study

Larger Data Set:

- Experiment repeated for data from 4450 devices with very similar results

Cost-driven compaction (TVLSI'09 – in press):

- Given appropriate test cost information (test time per performance, cost per test configuration, groups of performances sharing test configuration, per second cost of non-RF vs RF ATE, etc.) select subsets of performances that minimize cost instead of cardinality

Guard-banding:

- Assess the effectiveness of our guard-banding method in enabling exploration of test-cost / test quality trade-off

Current Research Activities (TI/IBM)

Dealing with process variations, shifts, and drifts:

- Investigate whether process variations over the lifetime of production affect the accuracy of the models
- Devise a method for adapting to process variations (periodic, event-driven, or monitoring-based retraining)

Dealing with variations across ATE and across sites:

- Investigate whether variations across sites of an ATE or across ATE affect the accuracy of the models
- Devise a method for adapting to ATE/site variations (training per site or ATE, using calibration filters, etc.)

Measurement correlations:

- Use correlations between PCM, wafer sort, and final test measurements to support adaptive test

What's Next? Stand-alone Built-in Self-Test

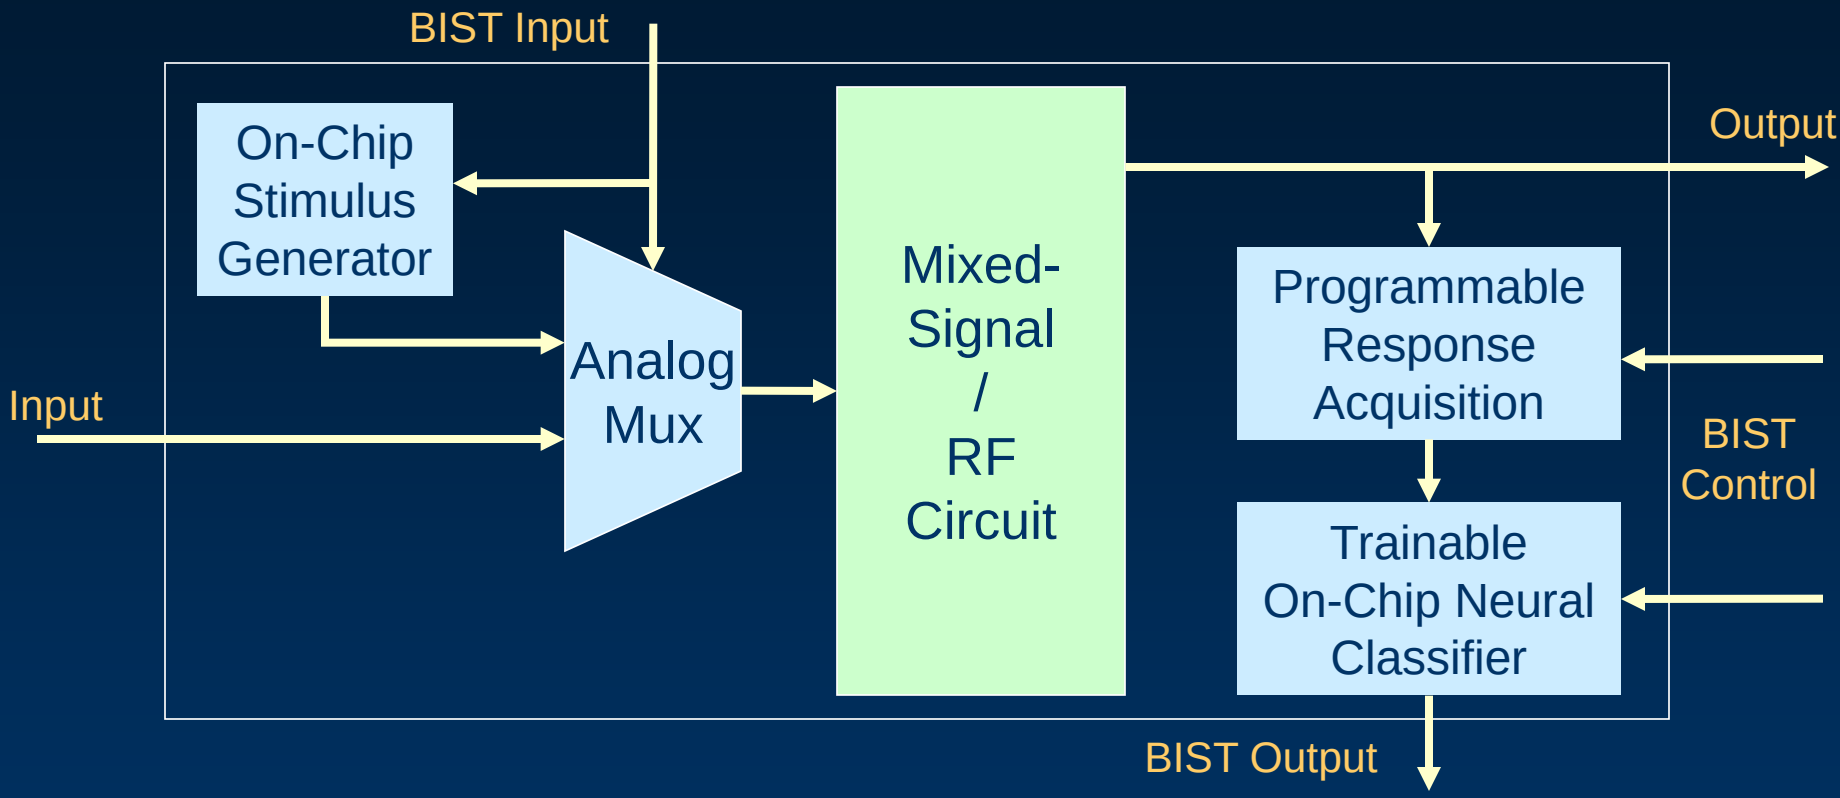

A stand-alone BIST method for mixed-signal/RF circuits:

- On-chip generation of simple test stimulus
- On-chip acquisition of simple measurements
- On-chip implementation of trainable neural classifier (w. floating gates)

Summary

Contribution:

- A non-linear neural classifier supporting a novel paradigm for testing analog/RF circuits
  - Achieves significant reduction in test cost without sacrificing accuracy of specification test
  - Enables exploration of trade-off between test cost and test accuracy via guard-bands

Applications:

- Disruptive: Machine learning-based test
- Non-disruptive: Specification test compaction
- Futuristic: Stand-alone Built-in Self-Test

Knob-Tuning for Performance Calibration

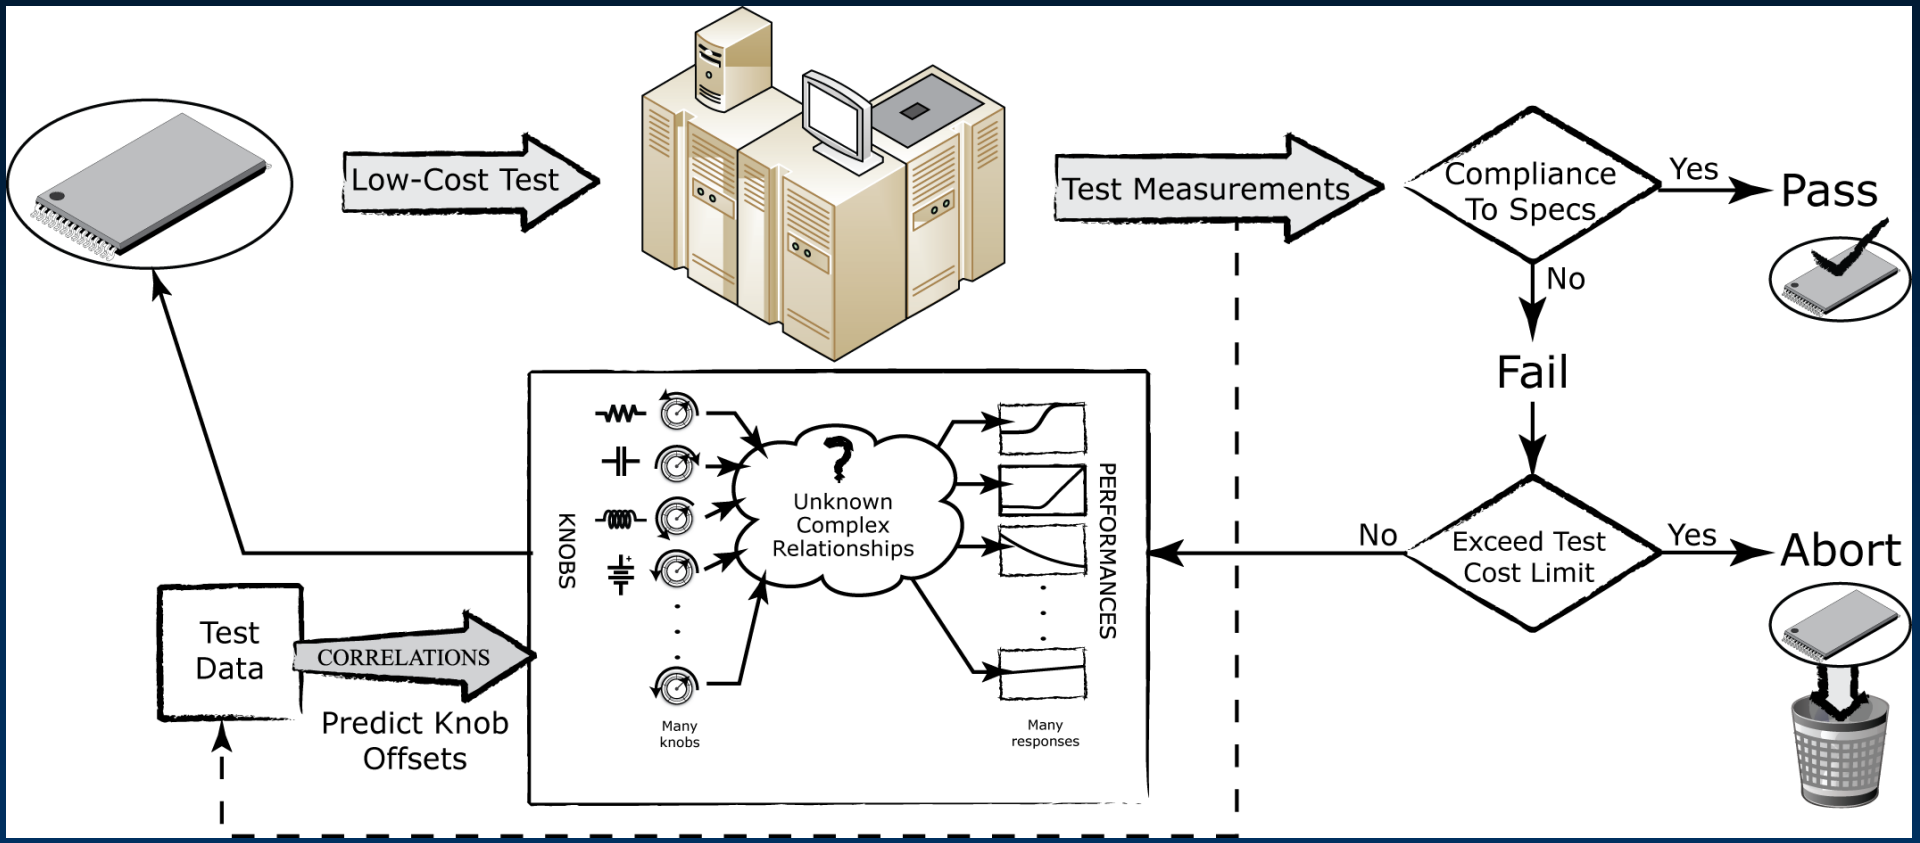

"Healing" of failing chips:

- Low-cost machine-learning based testing
- Iterative, correlation-based knob tuning

Correlation Mining for Yield Improvement

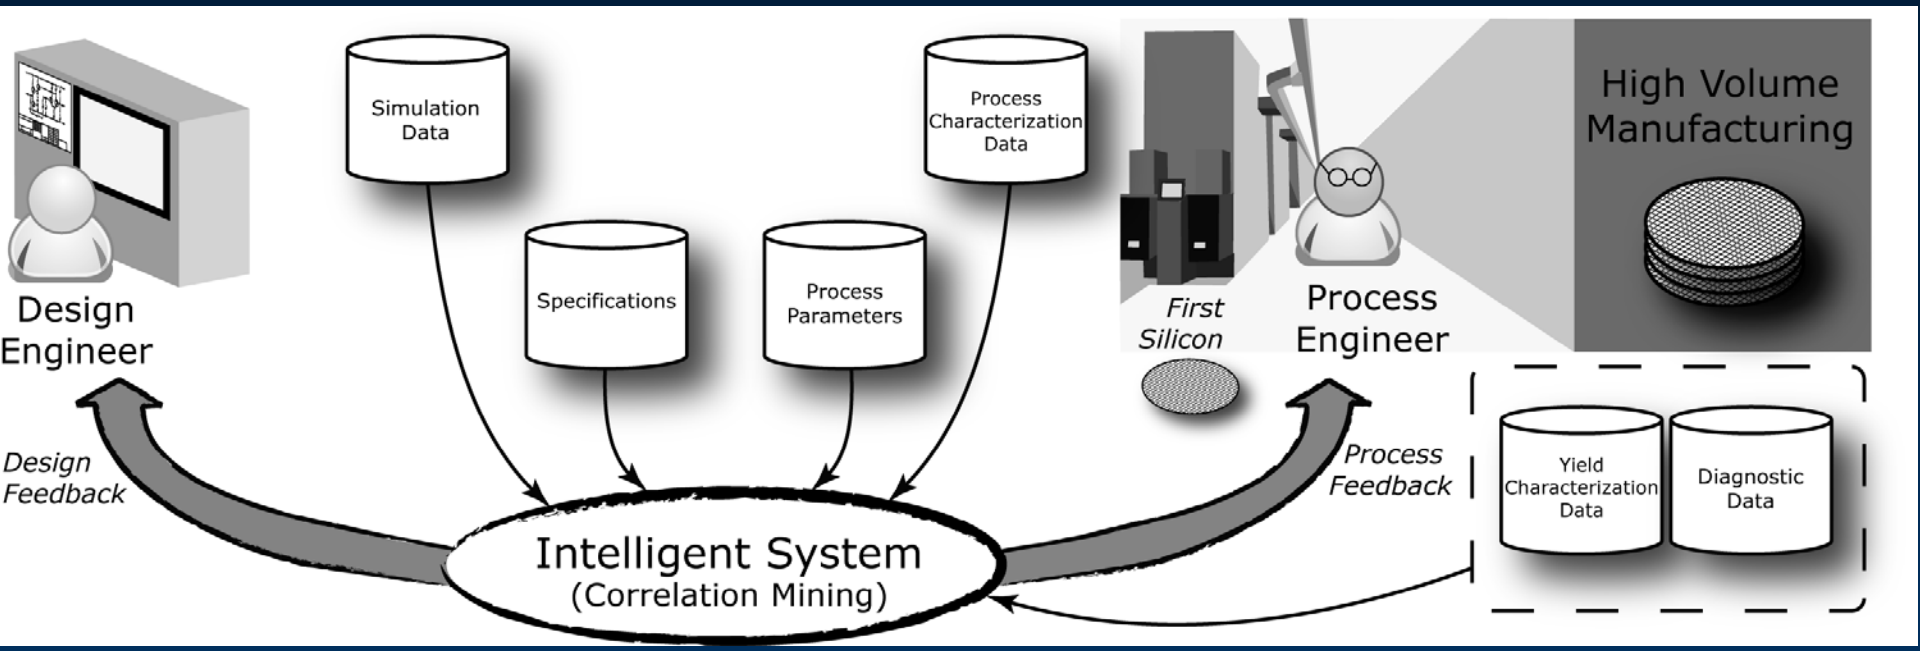

Improve yield by:

- Guiding modifications in potential design re-spin
- Assisting in fabrication process quality control

More Information

Publications:

- H-G. D. Stratigopoulos, Y. Makris, “Constructive Derivation of Analog Specification Test Criteria,” *Proceedings of the IEEE VLSI Test Symposium (VTS)*, pp. 245-251, 2005
- H-G. D. Stratigopoulos, Y. Makris, “Non-Linear Decision Boundaries for Testing Analog Circuits,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (T.CAD)*, pp. 1360-1373, Nov. 2005
- H-G. D. Stratigopoulos, Y. Makris, “Bridging the Accuracy of Functional and Machine-Learning-Based Mixed-Signal Testing,” *Proceedings of the IEEE VLSI Test Symposium (VTS)*, pp. 406-411, 2006
- H-G. D. Stratigopoulos, P. Drineas, M. Slamani, Y. Makris, “Non-RF to RF Test Correlation Using Learning Machines: A Case Study,” *Proceedings of the IEEE VLSI Test Symposium (VTS)*, pp. 9-14, 2007
- H-G. D. Stratigopoulos, Y. Makris, “*Error Moderation in Low-Cost Machine Learning-Based Analog/RF Testing* ,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (T.CAD)*, pp. 339-351, Feb. 2008

Contact:

- E-mail: yiorgos.makris@yale.edu

Questions?