



HAL
open science

Explainable AI for Embedded Systems Design: A Case Study of Static Redundant NVM Memory Write Prediction

Abdoulaye Gamatié, Yuyang Wang

► **To cite this version:**

Abdoulaye Gamatié, Yuyang Wang. Explainable AI for Embedded Systems Design: A Case Study of Static Redundant NVM Memory Write Prediction. 2024. lirmm-04493175

HAL Id: lirmm-04493175

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-04493175>

Preprint submitted on 6 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

EXPLAINABLE AI FOR EMBEDDED SYSTEMS DESIGN: A CASE STUDY OF STATIC REDUNDANT NVM MEMORY WRITE PREDICTION

Abdoulaye Gamatié, Yuyang Wang
LIRMM, Univ. Montpellier - CNRS
Montpellier, France
{firstname.lastname}@lirmm.fr

ABSTRACT

This paper investigates the application of eXplainable Artificial Intelligence (XAI) in the design of embedded systems using machine learning (ML). As a case study, it addresses the challenging problem of static silent store prediction. This involves identifying redundant memory writes based only on static program features. Eliminating such stores enhances performance and energy efficiency by reducing memory access and bus traffic, especially in the presence of emerging non-volatile memory technologies.

To achieve this, we propose a methodology consisting of: 1) the development of relevant ML models for explaining silent store prediction, and 2) the application of XAI to explain these models. We employ two state-of-the-art model-agnostic XAI methods to analyze the causes of silent stores. Through the case study, we evaluate the effectiveness of the methods. We find that these methods provide explanations for silent store predictions, which are consistent with known causes of silent store occurrences from previous studies. Typically, this allows us to confirm the prevalence of silent stores in operations that write the zero constant into memory, or the absence of silent stores in operations involving loop induction variables. This suggests the potential relevance of XAI in analyzing ML models' decision in embedded system design. From the case study, we share some valuable insights and pitfalls we encountered. More generally, this study aims to lay the groundwork for future research in the emerging field of XAI for embedded system design.

Keywords Explainable AI · machine learning · embedded systems · silent stores

1 Introduction

Machine learning (ML) is used in the design and optimization of embedded systems and Systems-on-Chip (SoCs) [1–4]. Examples include identifying memory access patterns for prefetching [5] and optimizing branch prediction [6]. ML techniques often handle complex, nonlinear relationships within a system. Complex ML models are difficult to understand or explain.

To tackle this challenge, *eXplainable Artificial Intelligence* (XAI) [7] has emerged, offering insights into model decision-making, and shedding light on potential biases and errors. The application of XAI to embedded system design aims to *enable designers to optimize nonfunctional properties such as power and energy consumption, as well as interpret and enhance the functional behavior of the system.*

Let us consider the design of an IoT device deployed in a smart home environment, where energy efficiency is critical. Data-driven modeling of the system design could be analyzed with XAI to examine how system components or functions contribute to overall energy consumption. Those insights could be used to redesign communication protocols or sensor sampling strategies during low activity periods. This would minimize energy consumption and reduce environmental impact. XAI-based iterative design of embedded systems for predicting equipment failures in industrial settings is another example. After initial deployment, users might find it challenging to trust the predictions due to occasional

unexpected outcomes. XAI can help uncover that some sensor readings, while critical, are not well understood by users. On the basis of this insight, designers can iterate the model by adding new sensor data or adjusting features to better match users' interpretation of equipment health.

1.1 Considered problem

Nowadays, embedded system designs using ML models often ignore the rationale behind the model decisions. Addressing this typically relies on costly *ad hoc* data analyses, as in *static silent store prediction* [8]. This nontrivial problem focuses on improving program execution by predicting, at compile time (hence the term *static*), which memory stores are likely to write the same value to a given location (hence the term *silent*).

Reducing silent stores helps mitigate the higher cost of writes compared to reads in emerging *non-volatile memories* (NVMs), thus reducing latency and energy consumption [8, 9]. During recent years, emerging NVMs have been extensively explored to improve the energy efficiency of embedded system architectures [10, 11]. In terms of latency and energy consumption, NVM memory writes are known to be more costly than reads. It also decreases write-backs and data bus traffic in multiprocessor systems [12]. In the *vortex* database benchmark from the SPEC CPU95 suite, up to 67% of memory stores are silent, and potential write-back reductions can reach 81% [12].

Eliminating silent stores is based on the *store-verify* method, where a *store* operation is replaced by three steps: a *load*, a *comparison*, and the *original store* itself. This process involves reading the target memory content, comparing it with the new value to be written; if they match, the store is skipped; otherwise, it is executed. Note that the store-verify method is similar to existing approaches such as the differential write method, which also suggests first reading the previous data in memory cells, comparing it with the new data, and only modifying cells whose values are different [13]. Using such a method, silent stores can be efficiently prevented at the hardware microarchitecture level. Nevertheless, it is not easily portable compared to software-oriented methods [8, 9], which integrate the store-verify optimization into program codes, ultimately making them executable on any hardware architecture.

To minimize the performance overhead caused by *store-verify*, it should only be applied to stores most likely to be silent. Hence, accurate silent store prediction is crucial, necessitating powerful ML models capable of capturing the intricate, nonlinear relationships between static program features and silent store occurrences, i.e., the *possible causality between both*. In this work, we consider static silent store prediction to demonstrate how XAI can be used in embedded systems design.

In this work, we tackle the challenge of building and explaining machine learning (ML) models for silent store prediction based on static program features. We focus on the following questions:

- (RQ1): *How can we build relevant ML models for identifying silent store causes based on their predictions?*
- (RQ2): *How can we effectively apply two state-of-the-art XAI methods for explaining these models?*

To deal with the above questions, we rely on a previous study [8], which already showed that static silent store prediction is beneficial to the energy efficiency of embedded systems integrating emerging non-volatile memory (NVM) technologies. To achieve this, the authors proposed experimental data to demonstrate the relevance of designing an ML-based prediction model to steer the effective application of the store-verify method. They trained a perceptron on a dataset. Afterward, they showed that depending on the benchmark, the energy cost of memory access can be improved up to a certain degree. In the present work, we use the same data set. However, we will focus on how to train effective machine learning models and how to interpret silent store predictions resulting from those models.

By addressing the aforementioned two questions, we identify some preliminary valuable insights into the application of XAI methods to static silent store prediction. Through this study, we also aim to illustrate how these methods could be adopted in ML-based embedded system design to explain model decisions.

1.2 Our contribution

We propose a two-step methodology to answer both (RQ1) and (RQ2).

- For (RQ1), we address a rationale for constructing suitable ML models tailored to silent store analysis.
- Regarding (RQ2), we define a sound and easily deployable approach for applying two XAI methods: *SHapley Additive Explanation* (SHAP) [14] and *Anchors* [15].

To validate this methodology, we conduct a case study to explore static program features that might influence silent store predictions. Our preliminary findings suggest that the used XAI methods provide explanations for silent store

predictions in accordance with previous studies that have identified possible causes of silent stores [8, 16]. For instance, they establish the frequent occurrence of silent stores in operations that write the zero constant into memory, or the absence of silent stores in operations involving loop induction variables, i.e., variable increment with integer constants. This is a key step towards leveraging XAI methods and improving memory access costs. This study not only confirms the effectiveness of these methods, but also uncovers potential pitfalls that can obscure XAI-based analysis results. Specifically, if the training dataset is unbalanced, precautions should be taken to ensure adequate prediction models are constructed for explainability. Finally, we briefly discuss a few design targets related to embedded systems that could benefit from XAI.

1.3 Outline

The rest of this document is organized as follows: Section 2 discusses some background notions and related work; then, Section 3 summarizes the reasoning methodology adopted in this work; Section 4 presents an application of the methodology in a case study to explain silent store prediction; Section 5 discusses some general observations made during our study; finally, Section 6 gives some concluding remarks.

2 Background notions and Related work

XAI finds application in a wide range of tasks, including decision support, predictive maintenance, and anomaly detection [17]. Depending on the ML models used, explanations can be generated either early in the training process by using *ante hoc* XAI methods or during inference by using *post hoc* methods [7]. *Ante hoc* methods are suitable for transparent ML models (understandable by human experts), such as k-nearest neighbors, linear regression, and decision trees. On the other hand, *post hoc* XAI methods rely on easily interpretable surrogate models that mimic complex base models with unknown inference mechanisms, like deep neural networks. These methods can be either *model-agnostic*, applicable to any ML model, or *model-specific*. Moreover, explanations can be either *local* or *global*. A *local* explanation aims to clarify why a model makes a prediction for a particular data instance. A *global* explanation seeks to elucidate how the model generates predictions in general, drawing insights from its entire training dataset.

Despite its potential, there is very little (if any) application of XAI to the design of embedded systems [18, 19]. The primary focus in [18] is on executing XAI on hardware accelerators, particularly to address transparency concerns with ML models. The concept of "explainable hardware" is introduced in [19] as a means of achieving system-level explainability through XAI, although in the preliminary stages. Our approach aims to construct non-trivial ML models and explain their predictions based on static program features, addressing an original set of research questions and providing insights not yet covered.

Furthermore, existing research on silent store detection such as [9, 12, 16, 20] is predominantly based on dynamic analyses, involving the execution and profiling of computer programs. However, a distinctive approach, defined at compile-time, was introduced in [8]. It aims at predicting silent stores based on static program features. This is very challenging as silentness is identified based on the syntactic properties of a program. This approach does not address the explainability of the considered ML models and cannot explain the implicit reasoning behind the silentness prediction. In the present work, we deal with this issue by using model-agnostic *post hoc* XAI methods. These explanations might later serve to avoid silent stores in program construction or compilation.

3 Reasoning methodology

We present a methodology for explaining silent store prediction.

3.1 Step 1: silent store classification

We address a binary classification problem, i.e., predicting whether a store is silent or not. In the latter case, the store is referred to as *noisy*. We consider two common metrics¹ for evaluating binary classifiers: *precision* and *recall*. These are particularly important when dealing with imbalanced data sets, such as the one we borrowed from [8].

¹Note that other metrics could also be considered to evaluate adequate NN models. Indeed, beyond classification, the design of embedded systems can involve *regression* or *clustering* problems. For regression, *Root Mean Squared Error* (RMSE), *Mean Absolute Error* (MAE), and R2 score are the most commonly used metrics. The first two determine how accurate predicted values are, and their deviation threshold from actual values. The latter measures the percentage of correct predictions returned by a model. Clustering problems use *similarity score* metrics, for example, the *silhouette coefficient* that measures how well-defined and distinct clusters are.

We denote *true positive* (a silent store predicted as silent) and *false positive* (a noisy store predicted as silent) by TP and FP , respectively. Similarly, we denote *true negative* (noisy store predicted as noisy) and *false negative* (silent store predicted as noisy) by TN and FN , respectively. *Precision* quantifies how many of all predictions are true positives, while *recall* measures how many actual silent stores are correctly predicted, i.e.: $precision = TP/(TP + FP)$ and $recall = TP/(TP + FN)$. Increasing *precision* by lowering FP in its denominator will result in a decrease in *recall* as TP will grow by definition but not FN , and *vice versa*. Then, how to decide a suitable balance between both metrics w.r.t. model explainability?

Let us consider the issue of silent store elimination for NVMs [8]. NVMs are known for their significantly higher write costs, both in terms of latency and energy, compared with read costs [9]. Upon a false positive prediction, applying the *store-verify* method imposes an unnecessary *read-and-compare* execution overhead. Assuming the *compare* operation is negligible compared with reading, let us denote the cost of a read by C_r . Upon a false negative prediction, an opportunity to replace a store with a more efficient *read-and-compare* operation is missed. In essence, if C_s represents the cost of a store, failing to make this replacement results in missing a potential savings of $C_s - C_r$. Now, let us assume a prediction model is prone to making a fixed number m of mistaken predictions, i.e., $FP + FN$. If this model predicts a percentage p of false positives, compared to a perfect predictor, it incurs overhead as follows:

$$m \times p \times C_r \tag{1}$$

Likewise, the proportion of FN for the same model will incur the following overhead compared with a perfect predictor:

$$m \times (1 - p) \times (C_s - C_r) \tag{2}$$

The total overhead of the model compared with a perfect predictor is defined by the sum of (1) and (2), as follows:

$$m \times (C_s - C_r) - m \times (C_s - 2 \times C_r) \times p \tag{3}$$

If $C_s > (2 \times C_r)$, increasing p in expression (3) will result in lower overhead. According to NVM technology designs, C_s can be 1.02 to 75 times higher than C_r [9]. The corresponding prediction model will have lower precision and higher recall. It is suitable for implementing the *store-verify* as shown in [8].

Conversely, when analyzing the significance of static features in relation to silent store prediction, it is advisable to employ more precise ML models that minimize the percentage p of FP as much as possible. Such models must have *higher precision than recall*. As a result, in this work, we will train ML models taking this objective into account for explaining silent store predictions.

3.2 Step 2: application of selected XAI methods

Combining global and local model-agnostic XAI methods, we examine the impact of static program features on silent store prediction. Global explanations are validated using local methods. In addition, we look for expressive and human-interpretable methods.

A fast global method we mentioned is *Accumulated Local Effects* (ALE) [21], which measures the marginal impact of a single feature on predicted output. By varying the value of the feature while keeping the other features constant, perturbed data instances are generated. The differences in predictions between perturbed and original instances are computed to determine local effects of a feature. In static silent store prediction, feature interaction effects are crucial as they favor silentness [8]. ALE is not well-suited to this. So, we use the *SHapley Additive* Explanation (SHAP) method [14]. This method allows ML models to be interpreted in a unified manner, accounting for isolated feature effects and feature interactions. SHAP is more computationally expensive than ALE.

To validate SHAP explanations, we exploit a fast local explanation method. The very popular *Local Interpretable Model-agnostic Explanations* (LIME) method [22] and its CLIME variant [23] are typical candidates. LIME explains models by learning a linear decision boundary that approximates a model w.r.t. a perturbation space. In CLIME, the user can specify Boolean constraints on this space to generate more focused perturbations. When applied to the static silent problem, these methods provide similar explanations for supposedly significant and less significant features. Consequently, it is difficult to discern the real impact of the different features. This limitation is possibly due to the way the methods handle feature space perturbation. We select the Anchors method, which explains non-linear boundaries in ML model behavior [15].

3.2.1 SHAP method

Let us consider a reference set Φ of features. SHAP computes the *Shapley value* of a feature $\phi_i \in \Phi$, w.r.t. a *feature vector* $\phi = \langle \phi_1, \dots, \phi_k \rangle$ and a function f corresponding to a prediction model. This value is the average contribution of feature ϕ_i to predictions by f w.r.t. ϕ , defined by function Ψ_i [14]:

$$\Psi_i(f, \phi) = \sum_{z' \subseteq \phi'} \frac{|z'|!(|\Phi| - |z'| - 1)!}{|\Phi|!} [f_\phi(z') - f_\phi(z' \setminus \phi_i)] \quad (4)$$

where ϕ' denotes a simplified interpretable binary vector that maps to the original input vector ϕ using a mapping function h_ϕ . Typically, $\phi = h_\phi(\phi')$ transforms a binary input vector ϕ' into the original input feature space, where $\phi'_i = 1$ (or 0) in the binary vector indicates the presence (or absence) of feature ϕ_i in the mapped original vector. The variable z' belongs to the set $\{0, 1\}^{|\Phi|}$ where $|\Phi|$ is the number of features in Φ . The notation $|z'|$ expresses the number of 1's in the binary vector z' ; and $z' \subseteq \phi'$ represents the set of all vectors z' where the 1's are a subset of those present in ϕ' . Here, $f_\phi(z') = f(h_\phi(z'))$ and the expression $[f_\phi(z') - f_\phi(z' \setminus \phi_i)]$ defines the marginal contribution of feature ϕ_i to the prediction function f w.r.t. an input vector z' . Note that $z' \setminus \phi_i$ means the value at position i in z' is set to 0. When evaluating f on feature vectors with fewer elements than the original input vectors, the missing features are commonly assigned default values. In this case, we have chosen to use *randomly chosen values* from the missing features within the working dataset D as default values.

In this study, each feature ϕ_i represents a specific static property of a store instruction $l: p[i]=v$, where l, p, i and v respectively denote the *instruction label* in the program code, the *pointer to the uploaded location*, the *offset added to the pointer* when building the target store address, and the *value to upload*. Table 1 summarizes the considered feature categories, based on [8]. Fig. 1 sketches four static feature vectors associated with store instructions.

```

1 #define SIZE 5
2 int main() {
3   int values[SIZE] = {1,2,3,4,5}; // (Vin, S10, Smn, Pay, Es1, ...)
4   int sum = 0; // (Ozr, Vin, sz8, S10, Smn, ZER, Scm, ...)
5   double average;
6
7   for (int i = 0; i < SIZE; i++)
8     { sum += values[i]; } // (Vin, sz8, S11, Smn, Oic, ...)
9   average = (double)sum / SIZE; // (Vdb, S10, Smn, Scm, DIV ...)
10  ... }
```

Figure 1: Static feature vectors specified within in C comments (lines 3, 4, 8 & 9)

All the features considered in this study are binary in nature: 1 means a specific feature is satisfied by a store instruction, while 0 means it is not. In general, as previously observed in [8], static silent stores are characterized by specific combinations of features, as illustrated in Fig. 1. Noteworthy feature combinations include:

- *nullifier* due to null value updates, contains Ozr or ZER; and
- *static memory initializer*, due to null value assignment in static memory, combines *nullifiers* with Msc.

As for noisy stores, notable feature combinations include

- *induction* due to variable increment with integer constants, which often involves ADD, Oin, Oic and INT.

In SHAP, dependence plots are provided that display the relationship between two different features based on the model's predictions. Nonetheless, this is not sufficient to address the combined effects of more than two static program features on store silentness. To overcome this limitation, we compute the *combined SHAP value* for vectors ϕ^k of k features, by proceeding as follows: 1) we define a set of ϕ^k to analyze; 2) for each such ϕ^k , we extract from the working dataset D , all static store data instances d_j containing all features of ϕ^k ; 3) finally, for each d_j , we average the sum of the SHAP values of its features to the silent store prediction.

More formally, the above algorithm can be defined as follows, with f representing the silent store prediction model:

1. define a set $\{\phi^k = \langle \phi_1, \dots, \phi_k \rangle, k \geq 2\}$ of k -feature vectors, containing feature combinations of interest; there are C_n^k possible combinations for each k , where $n = |\Phi|$;
2. for each vector ϕ^k , extract from the working dataset D all sets $D_{|\phi^k}$ of n -feature vectors containing the components of ϕ^k , as follows: $D_{|\phi^k} = \{d_{j,j \in 1..n} \in D, s.t. \phi^k \subseteq d_j\}$;

Table 1: Static program feature categories, applicable to store instruction 1: $p[i]=v$.

Categories of features	Examples of features with their corresponding encoding between parenthesis
the type of variable v	floating point on 32 bits (Vfp), floating point on 64 bits (Vdb), integer (Vin), pointer (Vpt)
the size in bits of the type associated with v	1 bit (sz1), 2 to 8 bits (sz8)
the program statements (or <i>program slice</i>) that may affect the values of v at some point of interest	addition operation (ADD), an LLVM get element pointer instruction (GEP), zero constant (ZER), integer constant (INT), subtraction operation (SUB), division operation (DIV)
the last instruction that produced the value of v	increment (Oic), assignment of zero (Ozr), integer constant update (Oin), load instruction (O1d)
the type of pointer p	integer (Pin), C-like struct (Pst), array (Pay), SIMD vector (Pvc)
the location of the region pointed by p	static memory (Msc), stack (Msk), heap (Mhp)
the label l of the store instruction in a program	within the main function (Snn), outside a loop (S10), within a singly, doubly, or triply nested loop (S11, S12, S13), compulsory store instruction (Scm)
the offset i that is added to p when building the store addresses, i.e., $p+i$	i has a stride of size 1 (Es1), or 8 (Es8), i is created by some affine expression, e.g., $i = 2*b + c$ (Eaf)

3. for each $D_{|\phi^k}$, average the sum of its feature components' SHAP values as follows:

$$\frac{[\sum_{j \in 1..n} \Psi_1(f, d_j)] + \dots + [\sum_{j \in 1..n} \Psi_k(f, d_j)]}{n \times k} \quad (5)$$

To evaluate the *soundness* of formula (5) w.r.t. combined feature importance, we also compute the exact ratio of d_j instances actually labeled as silent within $D_{|\phi^k}$, as follows:

$$\frac{|\{d_{j,j \in 1..n} \in D_{|\phi^k} \text{ s.t. } \text{label}(d_j, D) = \text{silent}\}|}{|D_{|\phi^k}|} \quad (6)$$

The validity of formula (5) is substantiated if the results obtained from formulas (5) and (6) are correlated for each vector ϕ^k , i.e., vectors exhibiting high (or low) combined SHAP values should correspond to high (or low) ratios of silent static store instances d_j .

3.2.2 Anchors method

Anchors [15] is an XAI method that provides human-understandable local explanations for specific data instances. An explanation is a predicate (also referred to as "rule") that, when satisfied, results in the same output predicted by the trained model for a given data instance to be explained.

In Anchors, a user may specify the portion of similar instances yielding the same prediction output, called *precision*. When precision requirements need to be high, Anchors generates highly restrictive predicates, which are typically satisfied by only a *small subset* of the dataset, commonly referred to as *coverage*. In particular, Anchors is less computationally intensive than SHAP.

4 Case study

We summarize the model training process. Then, we apply our methodology to assess feature importance in silent store prediction.

4.1 Model training for silent store prediction

We use the silent store dataset made available by [8] for our analysis. This dataset comprises a total of 89K static silent stores, each initially associated with 127 static program features. These stores were collected through the profiling of 222 programs. For each static store, the original dataset provides the proportion of silent *dynamic stores*, i.e., the number of executed instances of this static store that were silent during data collection. We, therefore, label a static store as "silent" only if all its dynamic stores are silent. About 10% of static stores are always silent. We reduce correlated features based on *Pearson correlation score*. This leads to 76 out of the original 127 features. Note that after cleaning the dataset, identical 76-feature instances emerged with different silentness labels (even though their original 127-feature versions did not). About 52% of original static store instances fall into this category. For our experiments, we use *Google Colab* as our computing platform.

We trained a multi-layer neural network (NN) using the above dataset. The NN model consists of five dense layers spanning from the input layer to the output layer. These layers have 64, 32, 16, 8, and 1 neurons, respectively. The first four layers use a *relu* activation function, while the output layer uses *sigmoid*. The dataset is split into 80% and 20% respectively for training and testing. For simplicity, we deliberately omitted a validation data subset. Here, having such a subset does not change our ML model's performance. The learning rate is $3e^{-4}$. The selected optimizer is *Adam*. We define 3000 training epochs (with *early stopping*) and a batch size of 20,000. The best precision and recall values obtained are about 0.60 and 0.29 respectively. These hyperparameters provide us with the best learning outcome among all the evaluated hyperparameter configurations.

4.2 Feature importance analysis with SHAP

In practice, the evaluation of SHAP values is computationally intensive. To mitigate this issue, we use an *explainer* of SHAP Python library [24], which approximates SHAP values for deep learning models, instead of the SHAP *exact explainer*.

4.2.1 Single feature perspective

Fig. 2 shows a SHAP beeswarm plot defined over 10 most influential features, among the 76 features mentioned in Table 1. The importance of features is shown from top to bottom, ordered by mean absolute SHAP values. It provides valuable insight into the impact of these features on the output of trained NN. Each point in the plot represents a specific data instance. Red and blue colors respectively indicate whether a store data instance satisfies or not a feature. The horizontal axis shows the SHAP value range of the features.

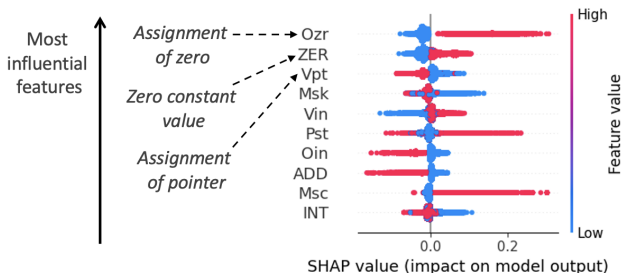


Figure 2: SHAP beeswarm plot over 10 most significant features w.r.t. silentness prediction by the trained NN model

The *nullifier* features OZR (i.e., stored value produced by an *assignment of zero*) and ZER (i.e., stored value is a *constant zero*) have the highest impact. When these features are satisfied by data instances, denoted by red points in the beeswarm plot, their SHAP values tend to increase. This improves the chance of predicting silent stores. As observed in [16], programmers frequently store zero as an explicit value or as a null pointer in their code. On the other hand, the *induction* features Oin (i.e., the last instruction that produces the stored value is an *integer constant update*) and ADD (that is, an *addition operation* contributes to stored value) exhibit the opposite trend. When satisfied by stores, the SHAP values tend to decrease, possibly contributing to store noisiness. Unlike previous work [8, 16], *loop-related features* indicating frequent store execution, do not highly contribute to store silentness in Fig. 2.

4.2.2 Combined feature perspective

While Fig. 2 highlights the impact of isolated features, combined feature effects can further increase silentness prediction chances as discussed below. We apply formulas (5) and (6) respectively to evaluate first the correlation between combined SHAP values and exact silentness ratios for various feature vectors. Then, we analyze feature combinations contributing the most to store silentness.

Given that the complexity of static silent store analysis grows with the number of features to be mined per instruction, smaller feature vectors are more cost-effective than longer vectors. In the sequel, we consider ϕ^k feature vectors with $k = 3, 4$, which are relevant for illustration. There are $C_{76}^3 = 70,300$ possible ϕ^3 vectors. Their enumeration on *Google Colab* takes around 30 minutes. For each ϕ^3 , we compute formulas (5) and (6). Enumerating all ϕ^4 vectors is computationally expensive, so we define a subset of ϕ^4 vectors by extending only ϕ^3 vectors with a combined SHAP value greater than 0.2. Fig. 3 summarizes the evaluation of our NN on three and four feature vectors respectively.

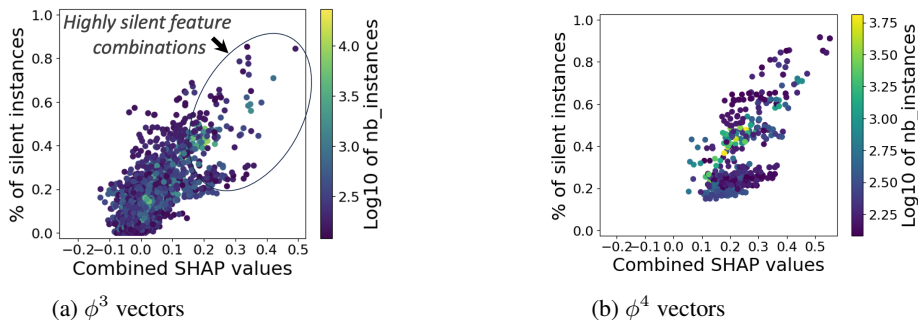


Figure 3: Ratio of silent stores vs combined SHAP values over 3- and 4-feature vectors for NN (each point denotes a vector)

The trend depicted in Fig. 3 reveals a reasonable correlation between the combined SHAP value (X-axis) and the exact silentness ratio (Y-axis) for ϕ^3 and ϕ^4 feature vectors. Each point on the plot denotes a k -feature vector. Its color represents the number (*log*. scale) of static store instances satisfying the vector features. This number varies from 120 to 25000 in Fig. 3. We deliberately omitted vector points associated with less than 120 static store instances as they are less representative. The correlation exhibited in Fig. 3 validates the soundness of combined SHAP values of features.

We notice that feature vectors with high combined SHAP values, i.e., between 0.38 and 0.55, often include the features `Ozr`, `Smn`, and `Msc`. Their respective contributions to the combined SHAP values of twelve ϕ^4 sample vectors are shown in Fig. 4. The *nullifier* feature `Ozr` stands out as the major contributor with an individual SHAP value above 0.21. In contrast, the other two features, `Smn` and `Msc`, contribute less than 0.17 individually. Interestingly, in the absence of the latter two features, the SHAP value of `Ozr` drops to 0.13.

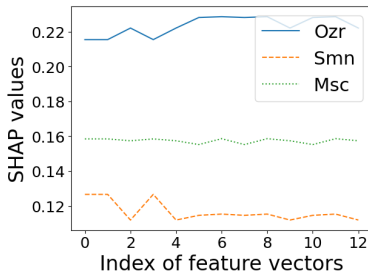


Figure 4: Feature contribution to silentness, over a few ϕ^4 vector samples in which `Ozr`, `Smn` and `Msc` always occur

Beyond the above three features, we also observe separately the features `ZER`, `Scm`, and `Vin` in ϕ^4 vectors that exhibit high combined SHAP values. The former has already been identified as an impactful *nullifier* feature. As for the other two features, they respectively describe *compulsory store instruction* in a program execution flow, and values `v` of *integer type* in a store instruction `l: p[i]=v`. Integer type, due to its smaller value range, contributes more to silent stores than floating point type. This was previously shown in [16]. This may explain the importance of `Vin` through its SHAP value.

For ϕ^4 feature vectors with medium combined SHAP values, i.e., between 0.22 and 0.36, the feature `Sz1` emerges as an additional notable contributor to silent stores. It often characterizes Boolean value updates, favoring silent stores as observed in [8].

4.3 Validating SHAP explanations with Anchors

Applying Anchors to the previous NN to generate local explanations requires a background dataset and the instance to explain. Here, we use the entire silent store dataset as the background dataset. The data instances to be explained are randomly chosen from those correctly predicted as silent. We impose that the generated "anchors" should have a precision of 0.95 at least.

The following predicate specifies an explanation generated by Anchors for a specific silent store instance:

$$(\text{Smn} > 0.00) \text{ AND } (\text{Ozr} > 0.00) \text{ AND } (\text{ZER} > 0.00) \text{ AND } (\text{Vin} > 0.00) \text{ AND } (\text{Scm} > 0.00) \text{ AND } (\text{INT} \leq 0.00)$$

This predicate essentially suggests static store instances satisfying (`Smn=1`, `Ozr=1`, `ZER=1`, `Vin=1`, `Scm=1`, `INT=0`), are likely to be classified as silent by the NN model. It includes *nullifier* features `Ozr` and `ZER`, combined with `Smn`, `Scm` and `Vin`. It denotes a null integer value assignment in a compulsory store executed in a main function. This is consistent with the results of the SHAP-based analysis. The combination of *nullifier* features with `Smn` and `Scm` favors store silentness [8].

4.4 Leveraging the resulting explanations to avoid silent stores

We observed from the SHAP and Anchor explanations provided in the previous sections that certain features are helpful indicators of silent stores. It is the case, for example, with *nullifier* features `Ozr` and `ZER`. Furthermore, we identified some feature combinations that are strongly associated with store silentness. In light of this knowledge, we can define a static program transformation that efficiently implements the store-verify method through a compiler. Given a program `P`, this process consists of the following steps:

1. analyze `P` and associate each store instruction in `P` with its corresponding static features;
2. apply the store-verify method to all store instructions with static features that strongly contribute to silent store occurrences.

By only replacing store instructions with a high silentness probability according to the obtained explanations, the above program transformation reduces the potential overhead associated with the store-verify method.

5 Overall discussion

We now discuss some insights and pitfalls from our study. The results of this study are useful for the future and the wider adoption of XAI in embedded system design.

5.1 Some insights and pitfalls

Insight 1: Although our ML models were moderate but acceptable in accuracy, the results obtained with the selected XAI methods are meaningful in general. We could not obtain very accurate ML models due to the inherent imbalance of the considered dataset. However, the derived explanations for static silent store causes were relevant and plausible in general.

User expertise is sometimes necessary to spot potential inconsistencies. For example, the seemingly minor influence of loop-related features indicated by XAI in our study raises questions about the validity of this observation w.r.t. existing literature [8, 16].

Insight 2: Training alternative ML models instead of neural networks does not necessarily improve the explanation of silent stores. Beyond NN, we also evaluated the Random Forest (RF) classification method, which is suitable for overcoming overfitting in decision tree training. RF harnesses the collective wisdom of weak decision tree learners to construct an accurate prediction model. Tree-based models inherently offer explainability [25], which makes RF an attractive candidate.

Consequently, we trained an RF model in which the resulting precision and recall values were 0.63 and 0.23, respectively (see Section 3.1). Fig. 5 shows a correlation between the silentness ratio and the combined SHAP value for ϕ^3 and ϕ^4 vectors.

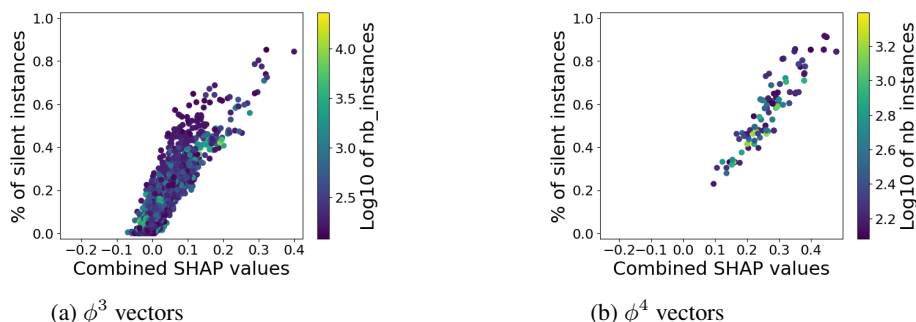


Figure 5: Ratio of silent stores vs combined SHAP values for RF

The number of ϕ^4 vectors obtained by extending ϕ^3 vectors with combined SHAP values above 0.2, is fewer than that for NN (see Fig. 3). However, the features that contribute to storage silentness are similar between the RF model and the NN model.

Insight 3: Choosing the appropriate trade-off between precision and recall in ML model training is crucial for an acceptable application of XAI to classification tasks, in presence of unbalanced dataset. Models with greater precision than recall are better suited for explainability given the unbalanced nature of our dataset (see Section 3.1). To support this claim, we evaluated the combined SHAP values for an NN with lower precision and greater recall. As depicted in Fig. 6, more ϕ^3 and ϕ^4 are predicted to be silent by this NN due to higher recall. However, the correlation with the ratio of silent vector instances is not as strong as in Fig. 3 due to a larger dispersion of vector points.

Given this precision/recall trade-off issue in static silent store prediction, does it mean that ML model explainability and code optimization goals are not compatible? This is a compelling question that deserves further investigation. More generally, we believe that the suitable trade-off also depends on the requirements of target problems. When false positives are critical in classification, higher precision is preferable. However, when false negatives are more critical, a higher recall is the better option.

Pitfall 1: Low combined SHAP values should not be systematically interpreted as noisy static store indications. Authors in [8] already identified the *induction* features `0ic` and `ADD` as strong indicators of noisiness. This is confirmed in our study through their low SHAP values, respectively -0.04 and -0.03 on average. However, this trend does not hold for combined SHAP values. Indeed, the working dataset D used in our study is imbalanced, with the majority of static store instances being noisy. To effectively distinguish actual noisy stores from those incorrectly predicted as noisy, more accurate ML models are necessary. For instance, in Fig. 3a, points with a negative combined SHAP values may

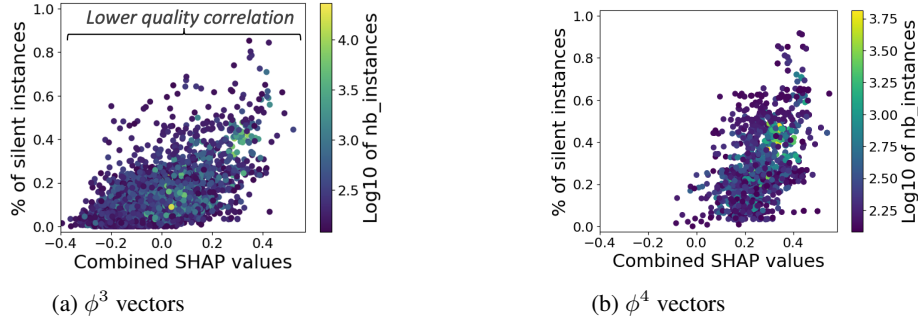


Figure 6: Ratio of silent stores vs combined SHAP values for NN, where precision is less than recall

have a silent store ratio between 0 and 0.6. This discrepancy likely results from inaccuracies in our NN, along with approximations in the employed SHAP *explainer*.

Pitfall 2: The outcomes produced by the Anchors method are not always easily exploitable. Beyond the above issue mostly related to SHAP explanations, Anchors predicates are sometimes needlessly complex, e.g., including several tens of irrelevant feature terms. From a human analysis standpoint, these lengthy predicates can make it challenging to extract the significance of highlighted feature properties. As an example, consider the following predicate of 31 feature terms, of which only two are satisfied, indicating the relevance of *nullifiers*:

(0zr > 0.00) AND (ZER > 0.00) AND (INT ≤ 0.00) AND (vpt ≤ 0.00) AND (SUB ≤ 0.00) AND ...

Such complex explanations can lead to confusion rather than insight. Consequently, a method like Anchors should be carefully manipulated by users for effective analysis.

Pitfall 3: Dataset quality concerning the store silentness problem affects the relevance of XAI-based explanations. Given the silent store data collection protocol used in [8], it appears that several identical feature vectors can land in different classes, i.e., noisy and silent. Around 60% of the dataset falls into this category. This means the dataset most likely does not accurately capture which features lead to silent or noisy stores. Among other factors, whether or not a store is silent also depends on the nature of the manipulated data. Consider the cumulative vector product as follows:

```
void prod(int* v, int* w, int* x, const int N) {
    for (int i = 0; i < N; ++i)
        x += v * w;
}
```

If either *v* or *w* contains zeros, one will have lots of silent stores; otherwise, the stores will be noisy. So, the same program can have very different behaviors w.r.t. store silentness, depending on its input data.

All features considered in [8] are syntactic, with the assumption that intuitively there is some relation between program syntax and semantics, although it is not very strong. For example, with the instruction `a = a + 1` the chance of a silent store is null. However, with the instruction `a = a + b*c`, if either *b* or *c* is zero, the store will be silent. Therefore, deciding whether the store is silent or not is not only affected by a syntactic analysis but also by a semantic analysis.

5.2 Towards XAI for embedded system design

Embedded system devices have unique characteristics, including dedicated functionality and various resource and behavioral constraints: limited processing power, memory, energy, size, real-time response, etc. In many existing studies, ML has been applied alone at both design and device operation [1, 2]. For instance, ML models can be customized for performance, temperature, or power prediction, ensuring devices' constraints are met. They can also focus on anomaly detection, enabling fault identification during device operation, and facilitating predictive maintenance. To make this possible, one needs to collect suitable datasets and train ML algorithms accurately based on adequate input features, as discussed in the following examples.

- To predict execution time, such features may include software characteristics (e.g., code complexity and size), execution platform specifications (e.g., processor architecture, clock frequency, and cache sizes), and hardware-related execution events (e.g., cache misses, branch mispredictions and memory accesses) [26]. For instance, code complexity can be measured using the diversity of instructions, functions or tasks [27], while code size can be measured using the number of lines of code. Cache misses are a key source of execution time, while cache size is a critical factor in determining the number of cache misses during execution.

- Power consumption prediction [28] can be based on input characteristics such as system configurations (e.g., voltage levels and clock frequencies), system operating conditions (e.g., CPU load, memory usage and network activity), and system environmental factors (e.g., temperature and humidity). Higher voltage levels and clock frequencies, as well as heavier CPU loads, can typically lead to higher power consumption.
- Embedded system temperature prediction [29] can be achieved using sensor data acquired within an embedded system, environmental conditions (e.g., ambient temperature and humidity), and system operational parameters (e.g., CPU usage and workload intensity).
- Input features relevant to security prediction [30] include network data (e.g., incoming and outgoing traffic patterns and packet headers), system logs (e.g., authentication attempts and access patterns), and behavioral patterns (e.g., user interactions and any abnormalities within the system). For example, network data could be used to detect suspicious traffic patterns or authentication attempts in a system. Logs could help identify compromised user spaces or suspicious system activity.
- Lastly, in the context of reliability prediction [31], relevant input features include failure data (e.g., previous instances of system failures or malfunctions), some testing data (e.g., data collected during stress tests or fault injection experiments), and environmental conditions (e.g., temperature, humidity and vibration levels during operation). Failure data may contain information about the failure type and when it happened. Testing data might indicate the performance scores of a system during tests and the number of conducted tests.

The last two examples are potentially relevant to application fields including integrated decision-making processes for autonomous vehicles, medical diagnosis, or energy distribution in smart grids.

While the majority of existing work focuses on ML model precision, understanding the factors that mostly influence model decisions is not covered enough. Given the high diversity of input features considered in the above examples, it is very challenging to address this concern. XAI may fill this need by providing designers with a means to explain or interpret built ML models. It could contribute to make embedded systems design more transparent and understandable. Nevertheless, the successful combination of ML and XAI in embedded system design is not trivial. This is illustrated through the insights and pitfalls reported in our study, on the intricate problem of static silent store prediction. This preliminary study paves the way to a wider investigation of XAI for embedded system design.

6 Concluding remarks

We considered XAI to study innovative approaches to embedded system design. This was illustrated through a pragmatic methodology applied to the static silent store prediction issue. Eliminating such redundant stores improves performance and energy efficiency. Leveraging cutting-edge XAI methods like SHAP and Anchors, we identified influential features from ML models to uncover some root causes of silent stores, aligning with prior research [8] and [16]. For example, this allowed us to confirm the frequent occurrence of silent stores in operations that write the zero constant into memory, or the absence of them in operations involving loop induction variables. The usage of explainable model predictions for silent store identification and elimination offers a new avenue for optimizing compiler techniques and hardware architectural design.

While XAI holds promise, this preliminary study emphasized, however, the need for cautious application, highlighting both valuable lessons and potential pitfalls to know. Finally, it lays the foundation for future research in the emerging area of XAI for embedded system design. To consolidate our current insights, further ML problems, e.g., including regressions, must be investigated in the future.

Acknowledgements

We thank Prof. Fernando Magno Quintao Pereira for many insightful comments and suggestions that contributed to improve this work.

References

- [1] Erika S. Alcorta, Andreas Gerstlauer, Chenhui Deng, Qi Sun, Zhiru Zhang, Ceyu Xu, Lisa Wu Wills, Daniela Sánchez Lopera, Wolfgang Ecker, Siddharth Garg, and Jiang Hu. Special session: Machine learning for embedded system design. In *2023 International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, pages 28–37, 2023.

- [2] Taiwo Samuel Ajani, Agbotiname Lucky Imoize, and Aderemi A Atayero. An overview of machine learning within embedded and mobile devices—optimizations and applications. *Sensors*, 21(13):4412, 2021.
- [3] Vijay Deep Bhatt, Wolfgang Ecker, Volkan Esen, Zhao Han, Daniela Sanchez Lopera, Rituj Patel, Lorenzo Servadei, Sahil Singla, Sven Wenzek, Vijaydeep Yadav, and Elena Zennaro. Soc design automation with ml – it’s time for research. In *2020 ACM/IEEE 2nd Workshop on Machine Learning for CAD (MLCAD)*, pages 35–36, 2020.
- [4] Martin Rapp, Hussam Amrouch, Yibo Lin, Bei Yu, David Z. Pan, Marilyn Wolf, and Jörg Henkel. Mlcad: A survey of research in machine learning for cad keynote paper. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 41(10):3162–3181, 2022.
- [5] Peter Braun and Heiner Litz. Understanding memory access patterns for prefetching. *International Workshop on AI-assisted Design for Architecture (AIDArc), held in conjunction with ISCA*, 2019.
- [6] Siavash Zangeneh Kamali. *Using convolutional neural networks to improve branch prediction*. PhD thesis, Univ. of Texas at Austin, 2022.
- [7] Alejandro Barredo Arrieta, Natalia Diaz-Rodriguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador Garcia, Sergio Gil-Lopez, Daniel Molina, Richard Benjamins, Raja Chatila, and Francisco Herrera. Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information Fusion*, 58:82–115, 2020.
- [8] Fernando Magno Quintão Pereira, Guilherme Vieira Leobas, and Abdoulaye Gamatié. Static prediction of silent stores. *ACM Trans. Archit. Code Optim.*, 15(4), nov 2018.
- [9] Rabab Bouziane, Erven Rohou, and Abdoulaye Gamatié. Compile-time silent-store elimination for energy efficiency: an analytic evaluation for non-volatile cache memory. In *Proceedings of the Rapido’18 Workshop on Rapid Simulation and Performance Evaluation: Methods and Tools*, RAPIDO ’18, New York, NY, USA, 2018. Association for Computing Machinery.
- [10] Sophiane Senni, Lionel Torres, Gilles Sassatelli, Abdoulaye Gamatié, and Bruno Mussard. Emerging non-volatile memory technologies exploration flow for processor architecture. In *2015 IEEE Computer Society Annual Symposium on VLSI*, pages 460–460, 2015.
- [11] Sophiane Senni, Lionel Torres, Gilles Sassatelli, Abdoulaye Gamatié, and Bruno Mussard. Exploring mram technologies for energy efficient systems-on-chip. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 6(3):279–292, 2016.
- [12] Kevin M. Lepak and Mikko H. Lipasti. On the value locality of store instructions. In *Int’l Symp. on Comp. Arch. (ISCA)*, page 182–191, 2000.
- [13] Ping Zhou, Bo Zhao, Jun Yang, and Youtao Zhang. A durable and energy efficient main memory using phase change memory technology. In *Proceedings of the 36th Annual International Symposium on Computer Architecture*, ISCA ’09, page 14–23, New York, NY, USA, 2009. Association for Computing Machinery.
- [14] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [15] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Anchors: High-precision model-agnostic explanations. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), Apr. 2018.
- [16] G.B. Bell, K.M. Lepak, and M.H. Lipasti. Characterization of silent stores. In *Proceedings 2000 International Conference on Parallel Architectures and Compilation Techniques (Cat. No.PR00622)*, pages 133–144, 2000.
- [17] Mir Riyanul Islam, Mobyen Uddin Ahmed, Shaibal Barua, and Shahina Begum. A systematic review of explainable artificial intelligence in terms of different application domains and tasks. *Applied Sciences*, 12(3):1353, 2022.
- [18] Zhixin Pan and Prabhat Mishra. Hardware acceleration of explainable machine learning. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1127–1130, 2022.
- [19] Timo Speith, Julian Speith, Steffen Becker, Yixin Zou, Asia Biega, and Christof Paar. Expanding explainability: From explainable artificial intelligence to explainable hardware, 2023.
- [20] Rabab Bouziane, Erven Rohou, and Abdoulaye Gamatié. How could compile-time program analysis help leveraging emerging nvm features? In *2017 First International Conference on Embedded & Distributed Systems (EDiS)*, pages 1–6. IEEE, 2017.
- [21] Daniel W Apley and Jingyu Zhu. Visualizing the effects of predictor variables in black box supervised learning models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(4):1059–1086, 2020.

- [22] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, page 1135–1144, New York, NY, USA, 2016. Association for Computing Machinery.
- [23] Aditya A. Shrotri, Nina Narodytska, Alexey Ignatiev, Kuldeep S Meel, Joao Marques-Silva, and Moshe Y. Vardi. Constraint-Driven Explanations for Black-Box ML Models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(8):8304–8314, Jun. 2022.
- [24] SHAP Python documentation, 2023.
- [25] Nadia Burkart and Marco F Huber. A survey on the explainability of supervised machine learning. *Journal of Artificial Intell. Research*, 70:245–317, 2021.
- [26] Tadeu Nogueira C. Andrade, George Lima, Veronica Maria Cadena Lima, Slim Bem-Amor, Ismail Hawila, and Liliana Cucu-Grosjean. On the impact of hardware-related events on the execution of real-time programs. *Des. Autom. Embedded Syst.*, 27(4):275–302, dec 2024.
- [27] Abdoulaye Gamatié, Xin An, Ying Zhang, An Kang, and Gilles Sassatelli. Empirical model-based performance prediction for application mapping on multicore architectures. *Journal of Systems Architecture*, 98:1–16, 2019.
- [28] Ajay Krishna Ananda Kumar and Andreas Gerstlauer. Learning-based cpu power modeling. In *2019 ACM/IEEE 1st Workshop on Machine Learning for CAD (MLCAD)*, pages 1–6, 2019.
- [29] Selma Rahman, Mattias Olausson, Carlo Vitucci, and Ioannis Avgouleas. Ambient temperature prediction for embedded systems using machine learning. In Jan Kofron, Tiziana Margaria, and Cristina Seculeanu, editors, *Engineering of Computer-Based Systems*, pages 12–25, Cham, 2024. Springer Nature Switzerland.
- [30] Basel Halak. *Machine Learning for Embedded System Security*. Springer, 2022.
- [31] Oussama Djedidi, Mohand A. Djeziri, and Samir Benmoussa. Remaining useful life prediction in embedded systems using an online auto-updated machine learning based modeling. *Microelectronics Reliability*, 119:114071, 2021.