



Prioritizing configuration relevance via compiler-based refined feature ranking[☆]

Federico Bruzzone^{ID}, Walter Cazzola^{ID}*, Luca Favini^{ID}

Università degli Studi di Milano, Computer Science Department, Milan, Italy

ARTICLE INFO

Dataset link: <https://zenodo.org/records/21172650>

Keywords:

Highly configurable systems
Testing variable systems
Configuration prioritization via static analysis
Rust

ABSTRACT

Modern programming languages, most notably Rust, offer advanced linguistic constructs for building highly configurable software systems as aggregations of features—identified by a configuration. However, they pose substantial challenges for program analysis, optimization, and testing, as the combinatorial explosion of configurations often makes exhaustive exploration infeasible. In this manuscript, we present the first compiler-based method for prioritizing configurations. Our approach consists of four main steps: (1) extracting a tailored intermediate representation from the Rust compiler, (2) constructing two complementary graph-based data structures—a *feature dependency graph* capturing inter-feature code-level dependencies weighted by predicate structure, and an *atom dependency tree* encoding the extent of code governed by each feature—, (3) using centrality measures to rank features, and (4) refining the ranking by considering the extent of code they impact. A fixed number of most relevant configurations are generated based on the achieved feature ranking. The validity of the generated configurations is guaranteed by using a SAT solver that takes a representation of feature dependencies in conjunctive normal form. We formalized this approach and implemented it in a prototype, *RUSTyEx*, by instrumenting the Rust compiler. An empirical evaluation on higher-ranked open-source Rust projects shows that *RUSTyEx* efficiently generates user-specified sets of configurations within bounded resources, while ensuring soundness by construction. The results demonstrate that centrality-guided configuration prioritization enables practical exploration of large configuration spaces, and we discuss the implications and open challenges for future configuration-aware analysis and optimization.

1. Introduction

Premise. Highly configurable software systems¹ often aim to satisfy a wide range of requirements. These systems provide a set of features that can be combined in different ways to create a variety of products (Apel et al., 2013). Variability-rich software system development leverages principles from product line engineering, commonly referred to as *feature-oriented programming* (Prehofer, 1997).² Programming languages, such as C/C++ and Java, provide mechanisms to manage software variability via preprocessor directives, conditional compilation, and annotations. Rust (Matsakis and Klock, 2014) has also been designed to support the development of highly configurable software systems via its *attribute system*³—specifically, the *cfg* attribute.

Problem Statement. Software variability can be complex, making it challenging to reason about all possible configurations (Agh et al., 2024). This problem is further exacerbated by the fact that the number of possible configurations grows exponentially with the number of features, as demonstrated by Krueger (Krueger, 2006). For instance, the Linux kernel is a well-known highly configurable software system with a significant number of features (Sincero et al., 2007) that exploits the native support for variability in C/C++ and, more recently, Rust (The Linux Kernel Organization, 2021). Over the years, the kernel and similar systems have been extensively studied, with numerous works highlighting the challenges posed by large configuration spaces (Melo et al., 2016; Tartler et al., 2011; El-Sharkawy et al., 2017; Martin et al., 2021).

[☆] Editor: Heiko Koziolk.

* Corresponding author.

E-mail addresses: federico.bruzzzone@unimi.it (F. Bruzzone), cazzola@di.unimi.it (W. Cazzola), luca.favini@studenti.unimi.it (L. Favini).

¹ These systems are also known as *product families* or *software product lines* (SPLs) (Clements and Northrop, 2001)

² We use this term loosely as a collective label for variability-implementation paradigms, not exclusively for the FOP paradigm of Prehofer

³ Rust's attributes are a form of syntactic metadata that can be attached to various parts of the code. Two kind of attributes are supported: *outer* and *inner*. The former is used to annotate *item* declarations (such as functions, structs, and enums), *expression* and *statement*. The latter is used to annotate *modules* and *block expressions* (in particular cases). We will refer to the *items*, *expressions*, and *statements* as *terms*. See <https://doc.rust-lang.org/reference> for more details.

State of the Art. Exhaustive approaches become infeasible due to the combinatorial explosion of configurations (Apel et al., 2011). Therefore, a strategy is needed to reduce the number of configurations under analysis. Several approaches aim to achieve this goal. Sampling techniques (Patel et al., 2013; Al-Hajjaji et al., 2019; Lee and Hwang, 2019), combinatorial testing (Oster et al., 2010; Lochau et al., 2012), and in particular t-wise testing (Cohen et al., 2007; Perrouin et al., 2010; Oster et al., 2010) attempt to reduce the configuration space while ensuring coverage. Other works have proposed prioritization methods (El-Sharkawy et al., 2019; Sánchez et al., 2014; Parejo et al., 2016; El-Sharkawy et al., 2017), using criteria such as similarity (Al-Hajjaji et al., 2019), non-functional properties (Sánchez et al., 2017), or centrality-based measures (Mohammed et al., 2024; Peng et al., 2015; Levasseur and Badri, 2024; Bagheri et al., 2010). As reported by Classen et al. (2013), many configurable systems are safety-critical, which makes prioritization even more important: the goal is to detect faults as early as possible while reducing the number of tests.

Limitations of Existing Approaches. Despite these efforts, the problem of finding a proper prioritization criterion remains open. Existing approaches 1. are predominantly based on feature models (Benavides et al., 2010; Pohl et al., 2011) and, while some recent work considers source artifacts or presence conditions (Hentze et al., 2022; Krieter et al., 2022), none combines code-extent weighting with centrality-based ranking in a compiler-integrated pipeline, and 2. do not account for the extent of the code affected by each feature. Consequently, configurations that are critical to the system's behavior may be omitted. Moreover, the need for principled configuration prioritization extends well beyond testing. Highly configurable systems must also address:

1. *compiler optimizations and performance analysis*, where either only a subset of variants can be feasibly evaluated, or configuration relevance can guide optimization strategies (Velez et al., 2020; Siegmund et al., 2015; Tartler et al., 2014; Velez et al., 2021),
2. *program comprehension and debugging*, where developers must reason about representative or critical configurations (Halin et al., 2019; Velez et al., 2022),
3. *variability management and refactoring*, where structural changes should be guided by the actual impact of features in the code (Liebig et al., 2015, 2017), and
4. *regression analysis*, where prioritization helps identify which configurations are most likely to reveal behavioral differences after evolution (Qu et al., 2008; Souto and d'Amorim, 2018).

In all these contexts, treating features uniformly or relying on stochastic heuristics may be insufficient, particularly in safety-critical and performance-sensitive settings where specific feature interactions must be systematically explored (Classen et al., 2013).

Proposal. In this paper, we propose, to the best of our knowledge, the first general method for prioritizing the relevance of configurations via a compiler-based refined ranking of features. To this end, we present RUSTYEx, a fully automated tool designed to identify the *most relevant configurations* in highly configurable Rust software systems.⁴ We define these configurations as those containing the most relevant features. A feature's relevance is measured by its centrality in the *feature dependency graph* and the extent of the code it impacts. As shown in Fig. 1, we instrument the Rust compiler by performing an interprocedural static analysis (Olender and Osterweil, 1990, 1992)⁵ to extract the *unified intermediate representation* (UIR). The UIR is obtained by removing irrelevant information from the Rust *abstract syntax tree* (AST) and by

creating the *atoms*.⁶ RUSTYEx performs static analysis on the UIR to extract two main data structures: the *dependency graph* (Mazurkiewicz, 1995) and the *polytree* (Dasgupta, 1999). The former is a weighted directed graph, dubbed *feature dependency graph* (see ② and ③ in Fig. 2), that represents the dependencies between the features. The latter is an induced subgraph of the UIR, dubbed *atom dependency tree* (see ④ and ⑤ in Fig. 2). This structure captures the *lexical scope*-based (Cooper and Torczon, 2022) dependencies between atoms, enriched with the extent of the code they affect. To address the previously mentioned limitations, we propose a method that combines both structures to identify the most relevant configurations. This method leverages centrality measures as structural metrics (Fenton, 1991; McCabe, 1976; Newman, 2010), and includes graph transformations into propositional and conjunctive normal form (CNF) (Classen et al., 2013) formulas, along with techniques for refining the ranking of features. RUSTYEx relies on a SAT solver to determine configurations that satisfy the CNF formula based on the refined feature ranking. This ensures that not only is the number of configurations reduced, but also that the most relevant configurations are prioritized. We validate our approach on higher-ranked open-source Rust projects by running RUSTYEx with a fixed number of configurations to generate. To provide a comprehensive evaluation, we report detailed metrics on the proposed structures for each project. Soundness is ensured by construction, guaranteeing that all generated configurations are valid.

Contributions. The main contributions of this paper are:

- the first general method for prioritizing configurations in highly configurable software systems via a compiler-based refined ranking of features,
- a formalization of the approach, from the extraction of the UIR to the algorithms for building the complementary data structures and prioritizing configurations,
- a detailed implementation of our approach in RUSTYEx, a fully automated tool for Rust software,
- an extensive evaluation of RUSTYEx on higher-ranked open-source Rust projects, demonstrating its effectiveness in identifying and prioritizing the most relevant configurations, and
- a formal proof of the soundness of our approach, ensuring that all generated configurations are valid.

Manuscript Structure. The remainder of the paper is organized as follows. Section 2 introduces the necessary background. Section 3 details the design of RUSTYEx, and Section 4 presents our evaluation. Section 7 discusses related work, and Section 8 concludes the paper.

2. Background

We provide background on Rust's configuration system and on the centrality measures used in this work.

2.1. The rust programming language

Rust is a system programming language that focuses on safety, speed, and concurrency (Matsakis and Klock, 2014). Rust supports software variability through its attribute system, particularly the *cfg* attribute, which allows conditional compilation based on specified configuration predicates. A *feature* in Rust is defined by all *terms* (i.e., items, expressions, and statements) annotated with the same *config name* in a *cfg* outer attribute. The inclusion of a feature in a product depends on the evaluation of a *configuration predicate* at compile time, e.g., `#[cfg(any(unix, windows))]`. Cargo—Rust's package manager—lets developers declare features using *config names* in the

⁴ Although we use Rust in this work, our approach to configuration prioritization is language-agnostic and applicable to any language with native variability support, such as C/C++ or Java.

⁵ Inter-procedural analysis is a static analysis technique that analyzes the control and data flow of a program across multiple functions.

⁶ From now on, we refer to an *atom* as a single *term* annotated with a given *configuration predicate*.

Cargo.toml file, combining them with logical operators into *configuration predicates*. Feature dependencies, i.e., *cross-tree constraints of feature models* (Kang et al., 1990; Seidl et al., 2016; Karataş et al., 2010), can also be specified in the same file. Unlike languages such as C/C++ that rely on preprocessor macros, Rust's *cfg* attributes are integrated into the compiler's type-checking and name-resolution pipeline, making them available for principled static analysis without macro expansion artifacts.

2.2. Centrality measures

Since the early days of social network analysis, *centrality measures* have been used to identify the key nodes in a network (Seeley, 1949; Bavelas, 1950; Katz, 1953). They have since become central to graph theory and network analysis (Borgatti, 2005; Das et al., 2018; Landherr et al., 2010). A *network* is typically modeled as a graph with n nodes, indexed by $i \in \{1, 2, \dots, n\}$, and represented as an adjacency matrix $A \in \mathbb{R}^{n \times n}$, where $A_{ij} \neq 0$ indicates an edge between nodes i and j , and $A_{ij} = 0$ otherwise. As Bloch et al. (2023) observe, not all centrality measures consider edge directions and weights, though these can be incorporated. Formally, a centrality measure is a function $\phi : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^n$, where $\phi_i(A)$ quantifies the relevance of node i in the network A . As a cardinal invariant, it enables ordinal ranking of nodes based on their scores ($\phi_i(A)$). Centrality measures are broadly categorized as: *geometry-based* and *spectral-based* (Boldi and Vigna, 2014). In the remainder of this section, we introduce the most commonly used centrality measures.

Geometric Centrality Measures. The *geometric* centrality measures assign relevance by using a function of distances—i.e., the number of nodes at each distance from a given node. *Degree centrality* counts the number of edges connected to a node: $d^-(i)$ for in-degree and $d^+(i)$ for out-degree. In a feature dependency graph, a feature with high in-degree is depended upon by many other features, suggesting it plays a central role in the configuration space. *Closeness centrality*, introduced by Bavelas (1950), defines a node's relevance as inversely proportional to the sum of its distances to all other nodes. Since the original formulation fails on disconnected graphs (due to infinite distances), a widely adopted variant is:

$$\text{Clo}_i(A) = \frac{1}{\sum_{j=1}^n d(i, j)}$$

where $d(i, j)$ denotes the shortest path between nodes i and j . In a feature dependency graph, a feature with high closeness centrality can quickly “reach” all others through dependency chains. *Harmonic centrality* (Marchiori and Latora, 2000) is defined as the sum of the reciprocals of the shortest path lengths between a node and all others:

$$\text{Har}_i(A) = \sum_{j=1}^n \frac{1}{d(i, j)}$$

Unlike closeness centrality, it is well-defined for disconnected graphs; for isolated nodes, the sum evaluates to 0 since $d(i, j) = \infty$ for all $j \neq i$. *Betweenness centrality* (Freeman, 1977) measures how often a node appears on the shortest paths between pairs of nodes:

$$\text{Bet}_i(A) = \sum_{s \neq i \neq t} \frac{\sigma_{st}(i)}{\sigma_{st}}$$

where σ_{st} is the number of shortest paths from s to t , and $\sigma_{st}(i)$ is the number of those paths passing through node i . A feature with high betweenness sits at structural crossroads: enabling or disabling it affects the reachability of many other feature combinations. Newman (2001) extended *closeness centrality* to weighted networks using Dijkstra's algorithm to compute shortest paths. Opsahl et al. (2010) further generalized shortest-path measures by combining edge weights and counts, making them applicable to both *closeness* and *betweenness* centrality.

Spectral Centrality Measures. The *spectral* centrality measures evaluate node relevance using the dominant left eigenvector of the graph's

adjacency matrix. *Eigenvector centrality* (Bonacich, 1972) assigns higher scores to nodes connected to other high-scoring nodes. It is defined as the eigenvector v associated with the largest eigenvalue λ of the adjacency matrix A :

$$\lambda v = Av$$

If A is a *stochastic matrix*, the dominant eigenvalue is 1. However, the *eigenvector centrality* performs poorly on disconnected graphs (Berman and Plemmons, 1979). In a feature dependency graph, a feature with high eigenvector centrality is not only itself important, but is depended upon by other important features—a property that degree centrality alone cannot capture. *Katz centrality* (Katz, 1953) addresses the disconnected-graph limitation by considering the number of paths of varying lengths that connect a node to all other nodes. It is defined as:

$$k = (I - \beta A)^{-1}$$

where $\mathbf{1}$ is a vector of ones, I is the identity matrix, and β is a damping factor satisfying $\beta < 1/\lambda$, with λ the dominant eigenvalue of A .

3. A deep dive into RustyEx

As introduced in Section 1, RustyEx is a fully automated tool that instruments Rust compiler to: 1. extract feature dependencies in Rust software, 2. assign feature weights based on their impact on the code, 3. apply centrality measures to rank features, and 4. prioritize configurations.

3.1. Process overview

RustyEx performs its analysis in three main phases: 1. *process setup*, 2. *dependency extraction*, and 3. *configuration generation*, as shown in Fig. 1.

Process Setup. Rust projects make extensive use of *cfg* attributes for conditional compilation, both for *feature gating*—enabling or disabling features via Cargo—and for purposes like platform-specific code selection (Anderson et al., 2016; Chen and Wu, 2022; Li et al., 2024). To ensure usability, RustyEx integrates seamlessly with the Rust toolchain. In this phase, it accepts a Rust project along with CLI options that customize the analysis, such as the number (N) of configurations to generate and the chosen centrality measure. The setup uses Cargo to resolve internal crate dependencies which are essential for the analysis.⁷ It then stores the parsed analysis configuration (centrality measure, number of desired configurations N , and crate list) in a shared in-memory state accessible to all compiler plugin instances spawned by Cargo, and configures the dispatch of source files to use the instrumented *rustc*.

Dependency Extraction. RustyEx instruments the Rust compiler to extract the *feature dependency graph* (see ②, and ② in Fig. 2) and the *atom dependency tree* (see ③, and ③ in Fig. 2) from the UIR derived from the AST. The details of these structures and their extraction are discussed in Section 3.2. This is integrated into the compiler by implementing the *rustc_driver_impl::Callback* trait, which provides four hooks: *config*, *after_crate_root_parsing*, *after_expansion*, and *after_analysis*.⁸ *rustc* performs the *eager* macro expansion⁹ after AST creation and before name resolution. Although, *after_crate_root_parsing* would be ideal for analysis before macro expansion, it lacks access to sub-modules that have not been resolved yet. Thus,

⁷ For performance, RustyEx analyzes only the target crate's source files and compiles external dependencies with the standard Rust compiler. This reduces computational overhead but may slightly impact accuracy (Venkatakeerthy et al., 2020).

⁸ See <https://doc.rust-lang.org/nightly/nightly-rustc/> for details.

⁹ Eager expansion expands macro arguments as early as possible, regardless of the macro invocation.




Figure 1. The phases of RustyEx (inspired from Li et al. (2021)).

we hook into `after_expansion`, which is invoked after `eager` macro expansion, name resolution, and AST validation. To avoid premature evaluation of `cfg` attributes, we register a custom file loader in the `config` callback. This loader renames `cfg` macros to delay their evaluation. In the `after_expansion` callback, we perform inter-procedural analysis on the AST to extract the UIR weighting its nodes with the *weighted* fixed-point algorithm shown in Algorithm 1 described in Section 3.2. These weights estimate each feature's impact on the code. **Configuration Generation.** Once the weighted feature dependency graph and the atom dependency tree are extracted, RustyEx uses them to identify the most relevant configurations. The weighted feature dependency graph serves two key tasks: 1. ranking features via *geometry-based* and/or *spectral-based* centrality measures (top-left box in the post-processing phase of Fig. 1), and 2. generating a corresponding propositional formula (top-right box in the post-processing phase of Fig. 1). These tasks are detailed in Section 3.3. The propositional formula is then converted to CNF. Next, the atom dependency tree refines the feature ranking (middle boxes in the post-processing phase of Fig. 1). Finally, a SAT solver selects the top N configurations that satisfy the CNF formula, prioritizing those with the highest refined ranking—i.e., the most relevant configurations (bottom box in the post-processing phase of Fig. 1).

3.2. From AST to UIR

Overview. In the AST, the `cfg` attribute nodes and their associated *terms* appear as separate child nodes under a common ancestor. Building the UIR involves two key steps: *unification* and *weighting*. The term UIR stems from the *unification* step, which merges each `cfg` attribute with its associated *term* into a single node called an *atom*. UIR nodes are algebraic data types (Lehmann and Smyth, 1981; Pierce, 2002; Turner, 1985; Kennedy and Russo, 2005; Bergstra and Tucker, 1995), specifically Σ -types—i.e., nodes can be either *atoms* or *relevant* plain AST nodes. Plain AST nodes are included in the UIR only if the *weighting* step assigns them non-zero relevance—e.g., generic bounds on parameters are usually excluded. For instance, in ❶ and ❷ of Fig. 2, the function `foo` and its `#[cfg(feature = "a")]` attribute on line 1 are unified into a single *atom* node, centered on the annotated item `foo`. Conversely, the body nodes of the function `bar` remain plain AST nodes in the UIR because they contribute to its weight, helping quantify how much of the code is influenced by the involved `cfg` features.

Formalization. We define the AST as a pair (N_{ast}, E_{ast}) , where N_{ast} and E_{ast} are the sets of nodes and edges, respectively. Let $N_{rel} \subseteq N_{ast}$ denote the set of *relevant* nodes. A node is considered *relevant* if it contributes to the semantics of a feature-dependent element. Specifically, `let` statements are relevant when they appear within the lexical scope of a `cfg`-annotated ancestor, as they contribute to the code extent of that scope; `let` statements in unconditional global scope with no `cfg` ancestor are excluded. Conversely, `extern crate` and

`use` declarations are never considered relevant, as they do not carry behavioral semantics for the purpose of feature ranking.

The set of *atoms* is defined as:

$$A = \{(p, t) \mid (ann(t) = p) \wedge (t \in T) \wedge (p \in P)\}$$

where:

1. P is the set of all `cfg` predicates, each defined as a recursive Σ -type with the following structure:

$$p = \begin{cases} \text{single}(f) & \text{for } f \in F \\ \text{not}(f) & \text{for } f \in F \\ \text{any}(p_1, p_2) & \text{for } p_1, p_2 \in P \\ \text{all}(p_1, p_2) & \text{for } p_1, p_2 \in P \end{cases}$$

2. $T \subseteq N_{ast}$ is the set of all *terms*, and
3. $ann : T \rightarrow P$ is the surjective function that maps terms to `cfg` predicates.

The UIR is an *enhanced* induced subgraph of the AST, where nodes are enriched with `cfg` predicates and edge directions are reversed. Formally, the UIR is defined as $\mathcal{U} = (N, E, w_N)$, where:

- $N = A \cup N_{rel}$ with $A \cap N_{rel} = \emptyset$, contains both *atoms* and *relevant* AST nodes, forming a Σ -type structure,
- $E = \{(i, j) \mid (i, j) \in E'\}$, where $E' \subseteq E_{ast}$ is the set of reversed edges derived from the AST, and
- $w_N : N \rightarrow \mathbb{N}^+$ assigns weights to nodes based on the type of *term*.

Since UIR nodes may be compound Σ -types (e.g., atoms), $E' \subseteq E_{ast}$ alone does not capture all necessary structure. To preserve connectivity, if $(i, j) \in E_{ast}$ and there exists $a \in A \subseteq N$ s.t. $i = \text{term}(a)$ or $j = \text{term}(a)$, we extend E' as:

$$E' \leftarrow E' \cup \{((ann(i), i), j) \vee (i, (ann(j), j)))\},$$

Here $\text{term} : A \rightarrow T$ extracts the term from an *atom*. This ensures that when a node is unified into an *atom*, its original AST edges are preserved in the UIR, as if the node remained *relevant* on its own.

Unification Algorithm. The *unification algorithm* performs a *depth-first* visit traversal of the AST to identify `cfg` attributes and their associated *terms*, merging them into *atoms*. RustyEx implements the visitor pattern (Gamma et al., 1995) via the `rustc_ast::visit::Visitor` trait. To track parent-child relationships during traversal, it maintains a *term stack*, where each visited term is pushed. Since `cfg` attributes always appear as leftmost children of terms, encountering one triggers the *unification* process. This involves:

1. extracting the configuration predicate from the `cfg` attribute,
2. parsing the predicate into a custom internal representation, dubbed `ComplexFeature`, and

3. popping the corresponding term from stack to *unify* it with the predicate into an *atom* node.

For example, the cfg attribute on line 3 of ① of Fig. 2 is parsed as

```
ComplexFeature::Any(vec[
  ComplexFeature::Simple(
    Feature {
      name: "b".to_string(),
      not: false
    }
  ),
  ComplexFeature::Simple(
    Feature {
      name: "c".to_string(),
      not: false
    }
  )
])
```

and unified with the term *bar* to form the corresponding *atom* node. Once a term is fully visited, it is popped from the stack and an UIR edge is added from the child to its parent. Unlike the AST, UIR edges are reversed to reflect lexical scope-based relationships, capturing dependency flow—key for applying graph centrality measures in the feature dependency graph (Section 3.3). The *bar* node in the atom dependency tree (② in Fig. 2) shows its unified features, highlighted by the ③ and ④ markers, emphasizing both unification and edge direction. **Weighting Algorithm.** Each node in the UIR is assigned a positive integers, i.e., elements of $\mathbb{N}^+$. During the *unification* phase, RUSTYEx assigns to every UIR node a *weight kind*, based solely on the kind of *term* it represents. These *weight kinds* determine how weights are computed in Algorithm 1 (lines 16–25), and include:

- *no weight*: nodes irrelevant to the analysis that receive no weight (e.g., **extern crate** declarations),
- *intrinsic weight*: nodes assigned a fixed weight of one (e.g., **let** statements),
- *children weight*: nodes whose weight is the sum of their children weights (e.g., *foo* in ① of Fig. 2), and
- *reference weight*: nodes that refer to other nodes in the UIR and inherit their weight (e.g., the *bar()* call in *qux* in ① of Fig. 2).

After *unification*, RUSTYEx runs a fixed-point algorithm (Algorithm 1) to compute the final weight function w_N over UIR nodes. In Rust, multiple *terms* can share the same identifier only if their cfg attributes are mutually exclusive. Let o_n denote the number of distinct *atoms* a node $n \in N$ appears in. The algorithm initializes:

- $m_w : N \rightarrow \mathbb{N}^{+o_n}$, a map storing all weights associated to the node identifier,
- an empty queue Q ,
- the root node $r = \text{root}(U')$, and
- the transposed UIR $U' = U'^T$, since the original UIR is built bottom-up and transposing it reestablishes parent-child relationships.

The *weighting* algorithm starts at the root node r , invoking the recursive `CALC_WEIGHT` function (line 17), which computes the weight of each node $n \in N$ by traversing its children in U' and applying the rules associated with its *weight kind*. Computed weights are stored both in the node and in the map m_w for reuse. For each node $n \in N$ of kind *reference*, such as function calls, the algorithm checks whether the referenced node already has a computed weight in m_w . If the weight is not yet available, the node is marked as *wait* and added to the queue Q . The fixed-point `RESOLVE_QUEUE` function (line 30) then processes Q , computing weights for the queued nodes until the queue is empty or a cycle is detected—e.g., due to direct or mutual recursion (von Oheimb, 1999; Magnusson and Hedin, 2007). A cycle is detected when Q remains

Algorithm 1 Calculate The Weight of the UIR Nodes

```
1:  $m_w \leftarrow \emptyset$ 
2:  $Q \leftarrow \emptyset$ 
3:  $r \leftarrow \text{root}(U')$ 
4:  $U' \leftarrow U'^T$  ▷ transpose
5: CALC_WEIGHT( $U'$ ,  $r$ ,  $Q$ ,  $m_w$ )
6: RESOLVE_QUEUE( $U'$ ,  $Q$ ,  $m_w$ )
7:  $U \leftarrow U'^T$  ▷ transpose back

8: function CALC_WEIGHT( $U'$ ,  $n$ ,  $Q$ ,  $m_w$ )
9:   for all  $adj \in \text{ADJ}(U', n)$  do ▷ adjacency nodes
10:     CALC_WEIGHT( $U'$ ,  $adj$ ,  $Q$ ,  $m_w$ )
11:   end for
12:    $ch_w \leftarrow 0$  ▷ children weight
13:   for all  $adj \in \text{ADJ}(U', n)$  do
14:     if  $adj.status = \text{WAIT}$  then ▷ check the status
15:        $n.status \leftarrow \text{WAIT}$ ; ret
16:     end if
17:      $ch_w \leftarrow ch_w + adj.weight$ 
18:   end for
19:   match  $n.weight\_kind$  with
20:     case No:  $n.weight \leftarrow 0$ 
21:     case INTRINSIC:  $n.weight \leftarrow 1 + ch_w$ 
22:     case CHILDREN:  $n.weight \leftarrow ch_w$ 
23:     case REFERENCE(called):
24:       if  $m_w[called] = \emptyset$  then ▷ no defs found
25:          $n.status \leftarrow \text{WAIT}$ 
26:          $Q \leftarrow Q \cup \{n\}$ 
27:         return
28:       else
29:          $t.weight \leftarrow \text{AVG}(m_w[called])$ 
30:       end if
31:   end match
32:    $n.status \leftarrow \text{WEIGHTED}$ 
33:    $m_w[n] \leftarrow m_w[n] \cup n.weight$ 
34: end function

35: function RESOLVE_QUEUE( $U'$ ,  $Q$ ,  $m_w$ )
36:    $S \leftarrow \emptyset$  ▷ set of seen nodes
37:   while  $Q \neq \emptyset$  do
38:      $c \leftarrow Q \downarrow$ 
39:      $Q \leftarrow Q - \{c\}$  ▷ dequeue
40:     if  $S \cap \{c, \text{LEN}(Q)\} \neq \emptyset$  then ▷ recovery mechanism
41:        $c.weight \leftarrow \text{DEF\_W}$  ▷ assign a default weight
42:       continue
43:     end if
44:      $S \leftarrow S \cup \{c, \text{LEN}(Q)\}$ 
45:     CALC_WEIGHT( $U'$ ,  $c$ ,  $Q$ ,  $m_w$ )
46:     if  $c.status = \text{WAIT}$  then
47:        $Q \leftarrow Q \cup \{c\}$ 
48:     end if
49:   end while
50: end function
```

unchanged between two complete iterations. To prevent infinite loops, the algorithm assigns a default fallback weight to nodes involved in cycles and continues processing the remaining queue.¹⁰ Finally, the UIR is transposed back to its original direction (undoing the earlier transformation), and the finalized weights are stored in w_N .

3.3. From the UIR to the feature dependency graph

Overview. Depicted in Fig. 2 (②, ③), the feature dependency graph is a weighted directed graph that represents dependencies between

¹⁰ Another approach has been proposed by instrumenting LLVM-IR (Venkateerthy et al., 2020), attempting to solve a system of linear equations that, by definition, could not admit solutions.




Figure 2. RUSTyEx process demonstrated on two scenarios.

software features. It is built by applying the *feature-extraction* algorithm to the UIR. This algorithm defines rules for assigning edge weights based on the *configuration predicates* of the *cfg* attributes in which the features appear. At a high level, the graph reflects lexical-scope-based relationships among features. Specifically, if f_i is used within the lexical scope of an *atom* annotated with f_j , then f_i depends on f_j and a directed edge is created from f_i to f_j . For instance, in line 3 of Fig. 2 ②, *feature* = "b" and *feature* = "c" appear in an *any configuration predicate*, within the lexical scope of an *atom* unified with the *feature* = "a" (line 1) and enclosing the *foo* function (line 2). According to the *feature-extraction* algorithm, two edges of weight 1 are created: from *feature* = "b" and from *feature* = "c" to *feature* = "a". In contrast, at line 7 of ②, *feature* = "c" is used in an *all configuration predicate* within the *global* lexical scope (denoted as **G**). Consequently, an edge of weight 1/2 is added from *feature* = "c" to **G**. As in the UIR, edge directions in the feature dependency graph are reversed compared to the original AST. The novel insight lies in preserving this natural edge direction: it allows us to interpret a feature's *reputation* (or relevance) based on the *state* (i.e., *configuration predicate*) of the features that depend on it. This enables the application of graph centrality measures on the feature dependency graph, where the edge weights are designed to support more accurate feature ranking.

Formalization. Until now, we referred to the feature dependency graph as a graph. However, during its construction, it is initially a *multigraph*, since a feature f_i may depend on another feature f_j in multiple parts of the code. Formally, this multigraph is denoted as $F = (F, D, w_R)$, where:

1. $F = \{(f_i, p_i) \mid f_i \in CN \wedge p_i \subseteq P\}$ is the set of nodes, where CN is the set of all *feature names* and p_i is the set of all *cfg predicates* in which f_i is used;
2. $D = (U, m)$ is the *multiset* of directed edges representing feature dependencies, where $U = \{(i, j) \mid i, j \in F\}$ is the set of edges and $m : U \rightarrow \mathbb{N}^+$ is the *multiplicity* function that counts how many times a dependency occurs; and
3. $w_R : R \rightarrow \mathbb{R}$, where

$$R = \{((i, j), k) \mid (i, j) \in U, 1 \leq k \leq m((i, j))\},$$

is the weight function that assigns a weight to each individual edge instance, indexed by the multiplicity index k , with weights determined by the *configuration predicates* in which the dependency arises.

The *feature-extraction* algorithm then squashes these multiple edges into single ones by aggregating their weights, effectively transforming the multigraph F into a simple graph $F' = (F, D', w'_R)$, where:

$$D' = \{(i, j) \mid (i, j) \in U\}$$

is a set of unique edges obtained by removing the multiplicities from D , and the aggregated weight function is defined as:

$$w'_R((i, j)) = \sum_{k=1}^{m((i, j))} w_R((i, j), k), \forall (i, j) \in D'.$$

Here, the sum aggregates the weights of all instances of the edge (i, j) , with $m(i, j)$ denoting the total number of such instances. Each individual

edge instance (i, j, k) contributes to the total weight assigned to the edge (i, j) in the final graph F' .

Feature-Extraction Algorithm. The *feature-extraction* algorithm builds the feature dependency graph F' from the UIR. It begins by iterating over the UIR atoms $A \subseteq N$. For each atom $a \in A$, it retrieves its parent node $\hat{a} \in \{x \mid (\hat{a}, x) \in E \wedge x \in A\}$. The existence of $\hat{a}$ and uniqueness are guaranteed for all $a \in A$ such that $a \neq \mathbf{G}$, since the UIR is an induced subgraph of the AST, and every node in a tree has a unique parent. This follows from the fact that the UIR is an *enhanced* induced subgraph of the AST. Concretely, $\hat{a}$ is the nearest ancestor of a that is an *atom*. The algorithm creates a directed edge from each feature $f \propto \text{pred}(a) = p_a$ to each feature $\hat{f} \propto \text{pred}(\hat{a}) = p_{\hat{a}}$, where $\propto$ is read as “involved in” and $\text{pred} : A \rightarrow P$ returns the configuration predicate that annotates each *atom*. The predicate sets of the *feature nodes* f and $\hat{f}$ are updated as:

$$(f, p) \leftarrow (f, p \wedge \text{pred}(a)) \text{ and } (\hat{f}, p') \leftarrow (\hat{f}, p' \wedge \text{pred}(\hat{a})).$$

The weights of the *multi-edges* are derived from the nested cfg predicate of a . For each feature node $\hat{f}$ and each $(f, w) \in \chi(p_a, 1)$, the algorithm creates a weighted edge from f to $\hat{f}$, using the recursively defined function χ :

$$\chi(p_a, w) = \begin{cases} [(f, w)], & \text{if } p_a \equiv \text{single}(f) \vee \text{not}(f) \\ \chi(p_1, 1) \oplus \chi(p_2, 1), & \text{if } p_a \equiv \text{any}(p_1, p_2) \\ \chi(p_1, d_{p_1}^w) \oplus \chi(p_2, d_{p_2}^w), & \text{if } p_a \equiv \text{all}(p_1, p_2) \end{cases}$$

where $d_p^w = (w \times |\{f \propto p\}|)^{-1}$ and the $\oplus$ infix operator concatenates two lists. For example, in ❶ of Fig. 2, the feature = “b” (line 3) appears in an any predicate nested under both feature = “a” (line 1) and foo (line 2), so an edge with weight 1 is added from feature = “b” to the feature = “a”. As discussed in Section 3.3, this weighting scheme prioritizes features in any predicates over those in all predicates, reflecting the fact that any predicates are less restrictive (only one feature needs to be enabled), whereas all predicates require all involved features to be enabled simultaneously. Finally, the graph is simplified by summing the weights of all edges with the same source and target nodes.

Centrality Measures. Once the feature dependency graph is built, RUSTYEX ranks features by relevance using a centrality measure specified via command-line argument (see Fig. 1, marked with a ● inside the *Centrality Vector* box). Since F is a disconnected, weighted, directed graph, some centrality measures are either undefined or behave unpredictably (Berman and Plemmons, 1979). For example, *eigenvector centrality* requires a connected graph (Berman and Plemmons, 1979) and both *closeness centrality* and *betweenness centrality* need adaptations to handle weighted directed graphs (Newman, 2001; Opsahl et al., 2010).¹¹ To ensure connectivity, we introduce a synthetic *patch* node (see ❷ and ❸ in Fig. 2). This node has a single incoming edge from $\mathbf{G}$ and outgoing edges to all other nodes, each with weight 1. This transformation guarantees that F is fully connected, allowing the following centrality measures to be applied without modification. Note that the patch node artificially elevates the centrality of features connected only to the global scope $\mathbf{G}$: in the worst case, globally-scoped features may receive centrality scores higher than their structural importance warrants. In practice, this effect is bounded, since globally-scoped features are already among the most broadly applicable. Users wishing to study this effect in isolation can disable the patch node via the -no-patch flag.

- *Newman Closeness Centrality* (Newman, 2001)
- *Opsahl Betweenness Centrality* (Opsahl et al., 2010)
- *Eigenvector Centrality* (Bonacich, 1972)
- *Katz Centrality* (Katz, 1953)

¹¹ RUSTYEX uses rustworkx (Treinish et al., 2022) to perform centrality measures.

After computing centrality scores, a vector $\vec{v}$ is produced where each entry v_i represents the centrality of feature f_i . The *patch* nodes are excluded from the vector, as they are not *extracted* features. This approach aligns with the intuition behind our graph construction: a feature’s reputation is defined by the features that depend on it. Centrality measures serve as a principled way to rank features by structural relevance.

Propositional Formula and CNF. From the feature dependency graph, RUSTYEX builds a *propositional formula* by iterating over each feature node $f_i \in F$. For each node, it generates a clause $\varphi_i = \ell(\text{pred}(f_i))$, where ℓ maps cfg predicates to *logical* expressions. It then adds an implication $f_i \Rightarrow f_j$ for each edge $(f_i, f_j) \in D'$, capturing feature dependencies. The overall formula φ is defined as:

$$\varphi = \bigwedge_{f_i \in F} \left(\varphi_i \wedge \bigwedge_{f_j \in \text{succ}(f_i)} (f_i \Rightarrow f_j) \right)$$

where $\text{succ}(f_i)$ is the set of successors of f_i in F' . Importantly, φ is *not* a global conjunction requiring all cfg predicates to be satisfied simultaneously. Each clause φ_i constrains f_i independently: if two code blocks carry conflicting predicates (e.g., $\#[\text{cfg}(a)]$ and $\#[\text{cfg}(\text{not}(a))]$), they produce two separate feature nodes f_a and $f_{\neg a}$ with their own independent clauses. A satisfying assignment α selects one consistent truth-value per feature, so conflicting predicates are handled by the SAT solver as disjoint alternatives—never as a single unsatisfiable conjunction. For example, in Fig. 2 ❹, activating feature = “a” produces a configuration that includes the foo function and, through the implications $f_b \Rightarrow f_a$ and $f_c \Rightarrow f_a$, also forces feature = “b” or feature = “c” to be enabled if they appear. Finally, φ is converted to CNF formula φ_{cnf} (Howson, 1997). Duplicate clauses may be removed before or after the CNF conversion.

3.4. From the UIR to the atom dependency tree

Overview. Fig. 2 (boxes ❸ and ❹) shows the atom dependency tree derived from the UIR. The name highlights its selective retention of UIR Σ -type nodes: only the *atom* variant is preserved, while *plain* AST nodes are removed. Specifically, only UIR nodes with cfg attributes are kept. For instance, in Fig. 2, box ❸, the gray qux node (representing the qux function from line 14 in Fig. 2, box ❶) is excluded because it lacks any cfg attribute. The atom dependency tree captures lexical scope-based dependencies among UIR *atoms*. Unlike the feature dependency graph, which flattens UIR structure, this tree preserves parent–child relationships between *atoms* and, by extension, between *features*. Each node is weighted by the amount of code affected by that atom’s cfg condition, reflecting the feature’s impact. To avoid losing information when *plain* AST nodes are discarded, the *atom dependency extraction* algorithm aggregates their weights into their nearest ancestor *atom* node (see foo node in Fig. 2, box ❸). Finally, the centrality vector $\vec{v}$ is refined using the structure of the atom dependency tree to get a better ranking of feature relevance.

Formalization. Given the UIR definition $\mathcal{U} = (N, E, w_N)$, we define the atom dependency tree as a directed graph $\mathcal{A} = (A, E_A, w_A)$, where:

- $A = N \setminus N_{\text{rel}}$ is the set of *atom* nodes from the UIR,
- $E_A = E \setminus \{(i, j) \mid (i, j) \in E \wedge (i \in N_{\text{rel}} \vee j \in N_{\text{rel}})\}$ is the set of edges only connecting *atom* nodes, and
- $w_A : A \rightarrow \mathbb{N}^+$ is the weight function that assigns to each atom a weight reflecting the amount of code affected by the cfg attributes in which its features appear.

The atom dependency tree ($\mathcal{A}$) is an induced subgraph of $\mathcal{U}$, containing all *atoms* and their dependencies, but excludes *plain* AST nodes. This holds trivially since $A \subseteq N$ and $E_A \subseteq E$.

Atom Dependency Extraction Algorithm. The *atom dependency extraction* algorithm builds the atom dependency tree from the UIR. It iterates over the UIR atoms $A \subseteq N$. For each atom $a \in A$, it identifies

its parent $\hat{a} \in \{x \mid (\hat{a}, x) \in E \wedge x \in A\}$. As before, the existence of $\hat{a}$ is trivial to prove. The algorithm then distinguishes two cases: 1. if $\text{pred}(a) \neq \emptyset$, it creates an edge from a to $\hat{a}$ and updates $\hat{a}$'s weight by adding a 's weight, 2. if $\text{pred}(a) = \emptyset$, it *does not* create an edge from a to $\hat{a}$ but still updates $\hat{a}$'s weight. For example, in Fig. 2, box ③, the qux node contributes to the weight of foo but does not create an edge. The parent weight update is performed as:

$$w_A(\hat{a}) \leftarrow w_A(\hat{a}) + w_N(a).$$

Refinement Algorithm. The *refinement algorithm* refines the centrality vector $\vec{v}$ using the atom dependency tree $\mathcal{A}$ (see Fig. 1). It iterates over each node $a \in A$ in $\mathcal{A}$, extracts every feature $f \in \text{pred}(a)$, and updates its centrality value in $\vec{v}$ as:

$$\vec{v}_f \leftarrow \vec{v}_f + \begin{cases} \frac{w_A(a)}{|\{f \in \text{pred}(a)\}|}, & \text{if } f \propto \text{single}(f) \vee f \propto \text{not}(f) \\ \frac{w_A(a)}{|\{f \in \text{pred}(a)\}|}, & \text{if } f \propto \text{any}(p_1, p_2) \\ \frac{w_A(a)}{|\{f \in \text{pred}(a)\}|}, & \text{if } f \propto \text{all}(p_1, p_2) \end{cases}$$

Here $f \propto p$, with $p \in P$, means that f appears in the p -variant of a 's cfg predicate, and $w_A(a) \in [0, 1]$ is the normalized weight of the node a . The normalization prevents the inclusion of the existing centrality values in $\vec{v}$. Although F captures dependencies between features, it may fall short of expressing their actual relevance. The atom dependency tree $\mathcal{A}$ refines centrality values by incorporating the *extent* of code affected by each feature, yielding $\vec{v}'$ a new permutation of the vector $\vec{v}$.

3.5. Configuration generation

Configuration generation is the final step of RustyEx. It produces the most relevant configurations based on the centrality vector $\vec{v}'$ and the CNF formula φ_{cnf} (see Fig. 1). While more advanced lexical-scope-based strategies could be considered, RustyEx uses SAT solvers—both incremental (e.g., MiniSat (Eén and Sörensson, 2003; Sörensson and Eén, 2005) and CaDiCaL (Biere et al., 2020)) and non-incremental (e.g., Kissat (Biere et al., 2020)). Given a fixed number K of configurations to generate, the algorithm iterates over the values in $\vec{v}'$. For each feature f_i , it updates the formula as:

$$\varphi_{\text{cnf}} \leftarrow \varphi_{\text{cnf}} \wedge f_i$$

Then, it uses a SAT solver to generate all the configurations satisfying φ_{cnf} . A SAT solver is then invoked to produce all satisfying configurations for the updated formula. If a satisfying configuration is found, it is added to the result set. This process repeats until K configurations are obtained, negating each newly found configuration to favor the discovery of new ones.

3.6. Correctness of configuration generation

Premises. To prove correctness, we show that: 1. the CNF encoding faithfully represents all configuration constraints, 2. the SAT solver returns only satisfying assignments, and 3. no invalid configuration is ever generated.

Construction of the CNF Formula. *Clause formation*—for each feature node $f_i \in F$, a clause $\varphi_i = \ell(\text{pred}(f_i))$ is created to enforce the logical interpretation of the feature's configuration predicate. That is, if a feature is enabled in some configuration, then its predicate must be satisfied. *Implication formation*—for every outgoing edge to f_j from f_i (i.e., for each f_j such that $(f_i, f_j) \in D'$), an implication $f_i \Rightarrow f_j$ is added to capture the dependency relations. This ensures that if feature f_i is active, then all its dependent f_j are also active. *Extension*—the formula φ is extended to include mandatory features specified in Cargo.toml, guaranteeing their presence in all generated configurations. *Equivalence Preservation*—since the CNF conversion uses standard techniques that preserve logical equivalence (Howson, 1997),

any assignment that satisfies the CNF formula φ_{cnf} also satisfies the original formula φ .

Correctness of the SAT Solution. Given a SAT solver sound and complete,¹² any assignment α returned for φ_{cnf} satisfies every clause in φ_{cnf} , and thus also in φ . This implies:

1. for each feature f_i with clause φ_i , its activation under α satisfies its configuration predicate: $\alpha \models \varphi_i \quad \forall f_i \in F$ —that is, α models the predicate;
2. for each dependency $f_i \Rightarrow f_j$, the implication holds under the assignment α if $\alpha(f_i) = \text{true} \Rightarrow \alpha(f_j) = \text{true}$;
3. all mandatory features from Cargo.toml are active in α : $\forall f_i \in F_{\text{mandatory}}, \alpha(f_i) = \text{true}$.

By definition, a configuration is valid if and only if:

- it satisfies all the configuration constraints including feature predicates and dependencies, and
- it includes all mandatory features.

Therefore, since any assignment α produced by the SAT solver satisfies all clauses in φ_{cnf} , every configuration generated by RustyEx is valid.

Reinforcement by Contradiction. Assume, for contradiction, that the SAT solver returns a configuration α' that is invalid—that is, it violates one or more constraints. Then there exists at least one clause in φ_{cnf} that is *not* satisfied by α' . By definition, a *satisfying assignment* satisfies all clauses in the formula, and the soundness of the SAT solver guarantees that α' is a satisfying assignment for φ_{cnf} . This contradiction implies that our assumption is false: α' cannot be invalid. Hence, all configurations produced by the SAT solver are valid.

3.7. Design rationale

We now justify the key design choices of RustyEx, discussing alternatives considered and their trade-offs.

Why a Feature Dependency Graph rather than a Feature Model?

A *feature model* (Kang et al., 1990) is a declarative artifact that must be explicitly authored and maintained alongside the software. Many modern Rust projects do not provide a formal feature model; instead, variability is encoded implicitly through cfg attributes and Cargo.toml dependencies. The *feature dependency graph* (FDG) is inferred automatically from the source code and build metadata, requiring no additional modeling effort. Moreover, the FDG captures code-level transitive dependencies that a feature model typically does not represent—e.g., the fact that enabling feature b within the scope of feature a creates a dependency that a high-level FM would need to be explicitly encoded.

Why an Atom Dependency Tree in addition to the FDG? The FDG records *which* features depend on each other, but not *how much code* each feature governs. Two features with identical FDG centrality may differ dramatically in their impact: one may gate a single statement while another gates an entire module. The *atom dependency tree* (ADT) fills this gap by tracking the weighted extent of code affected by each cfg-annotated atom. The refinement step (Section 3.4) adjusts centrality scores using these weights, yielding a ranking that is both structurally and quantitatively informed. A flat weight vector without the ADT topology was considered but discards the parent-child scoping information needed to correctly aggregate weights across nested cfg attributes.

Why Geometric and Spectral Centrality Measures? Simple degree centrality is insensitive to indirect influence: a feature with many transitive dependents may have low degree. We include both geometry-based measures (Newman closeness, Opsahl betweenness), which are robust to disconnected graphs and capture path-based influence, and

¹² A SAT solver is *sound* if it only returns assignments that satisfy the formula, and *complete* if it finds a solution whenever one exists.

spectral measures (eigenvector centrality, Katz centrality), which propagate relevance recursively through the graph. The choice of measure is exposed as a user parameter, because no single measure is universally optimal across graph topologies (Boldi and Vigna, 2014). Offering a configurable portfolio lets practitioners match the measure to the characteristics of their project's dependency graph.

Why SAT Solving for Configuration Generation? Alternatives considered include random-walk enumeration and BDD-based enumeration. Random walks cannot guarantee coverage of the highest-ranked features. BDD-based approaches can enumerate configurations exhaustively, but their memory usage grows exponentially with the number of variables (Apel et al., 2013). SAT solvers handle large CNF formulas efficiently in practice and can be guided by assumption literals to prioritize configurations that include the top-ranked features—exactly the behavior required. Both incremental solvers (MiniSat, CaDiCaL) and non-incremental solvers (Kissat) are supported, allowing users to trade-off between per-query overhead and overall throughput.

Why a Fixed Timeout per Crate? Analyzing each crate with an individual timeout (rather than a global project-level budget) ensures fair comparison across projects regardless of workspace size. A fixed per-crate timeout bounds the worst-case cost of any single analysis unit, enabling RUSTYEx to make progress on the remaining crates even when one crate is pathologically large or has missing dependencies.

3.8. Applications and clarifications

The generation of configurations in RUSTYEx is performed incrementally and in a *lazy* fashion. The tool ranks all features according to their refined centrality in $\vec{v}'$ and then generates only the top K configurations most likely to cover critical feature interactions. In practice, this means prioritizing configurations that include the highest-ranked features in $\vec{v}'$, starting with $\vec{v}'_0$, the most central feature.

As discussed earlier, this prioritization strategy supports efficient exploration of the configuration space by focusing on variants that are most representative and impactful in terms of feature relevance. By analyzing these configurations first, developers can uncover important interactions and behaviors in the most relevant scenarios, improving overall system understanding and validation. Moreover, this targeted approach significantly reduces the computational cost and resource consumption associated with exhaustive analysis of all configurations, making it a practical solution for large-scale, highly configurable systems. For instance, these prioritized configurations could also serve as candidates for targeted performance profiling or static analysis, ensuring that critical feature combinations are examined before less relevant ones.

This configuration-aware pipeline highlights the potential of centrality-guided heuristics in feature-oriented analysis of configurable systems and paves the way for future work in optimizing configuration sampling based on structural and semantic program properties. It suggests a broader paradigm where program structure and semantics guide systematic exploration of the configuration space, potentially benefiting tasks such as performance evaluation or security analysis (see Table 1).

4. Evaluation

4.1. Research questions

Our evaluation is organized around four research questions:

RQ1 (Soundness). Are all configurations generated by RUSTYEx valid with respect to the feature constraints of the analyzed project?

RQ2 (Efficiency and Scalability). Does RUSTYEx scale to large, real-world Rust projects in terms of execution time and memory usage?

RQ3 (Effectiveness). Do RUSTYEx-prioritized configurations outperform baselines in terms of fault-detection capability or code coverage?

RQ4 (Robustness). What is the failure rate of RUSTYEx across diverse Rust projects, and what are the root causes?

4.2. RQ1: Soundness

The soundness of RUSTYEx's configuration generation is guaranteed by construction, as formally proved in Section 3.6. Here we confirm the guarantee empirically: across all 1628 crates successfully analyzed in our study, no invalid configuration was produced. Every generated configuration was verified to satisfy the corresponding CNF formula φ_{cnf} and to include all mandatory features declared in Cargo.toml.

Answer to RQ1. RUSTYEx generates exclusively valid configurations, both by formal construction and by empirical observation across the full evaluation dataset.

4.3. RQ2: Efficiency and scalability

Experimental Setup. All experiments were conducted on a laptop equipped with an Intel Core i5-1135G7@2.40 GHz CPU and 16 GB of RAM. This choice highlights the lightweight nature of our approach: unlike some program analysis and variability-management tools that require dedicated servers or high-performance hardware (Kröher et al., 2018; Kenner et al., 2010), RUSTYEx can be executed on standard developer workstations. Most of the analyzed projects are organized as Cargo workspaces, which naturally decompose the system into multiple crates. For each crate, we executed RUSTYEx with a timeout of 10 min. While this bound is above typical every-commit CI feedback targets, it is appropriate for periodic integration checks such as nightly builds or pre-release gates (Hilton et al., 2016); a discussion of CI suitability appears in Section 5. Various metrics were collected and then aggregated at the project level, including counts of features and dependencies, statistics on UIR nodes and edges, and summaries of the feature dependency graph and the atom dependency tree. We also monitored peak memory usage and execution time. A replication package with all experimental data and scripts is available on Zenodo.¹³

Execution Time and Scalability. Fig. 3—highlighted in ●—shows the correlation between lines of code and execution time on a logarithmic scale. RUSTYEx successfully completed the analysis of about 93% of the projects, showing strong scalability across a wide range of sizes and domains. Execution time scaled smoothly with code size: for small to medium projects, the analysis often completed within seconds, while for the largest ones (e.g., swc at 648,247 LoC) execution time peaked at 33 min (2016 s; this data point lies above 10^3 s on Fig. 3 and appears at the top of the y-axis range). The average execution time per project was 333 s; the median was 163 s, which better represents the typical case as the mean is skewed by the two largest outliers (swc and diem). Individual crates required about 8 s on average. These results confirm that RUSTYEx is efficient in practice and robust enough to handle large modular codebases. Importantly, the logarithmic trend visible in the plot suggests that the analysis cost grows sub-linearly compared to project size. We note that the 33-minute peak is not suitable for every-commit CI latency; see Section 5 for a detailed discussion of practical adoption scenarios.

Memory Usage. Fig. 3—highlighted in ●—shows the correlation between lines of code and peak memory usage, also on a logarithmic scale. Memory consumption was generally modest: most projects required less than 1 GB of RAM, while only a handful of large codebases such as hyperswitch and swc reached a peak of 3.8 GB. The average peak usage was 885 MB per project, which fits comfortably within the

¹³ <https://doi.org/10.5281/zenodo.21172650>

Table 1
Results of the experiments conducted on 40 Rust projects. All columns are aggregated per project, summing the values of all crates within the workspace, except for *Peak Memory Usage* and *UIR Height*, which are the maximum values.

Project Name	GitHub stars	Crates.io downloads	Lines of code	Workspace members	Failed members	Dependency	Defined features	UIR nodes	UIR edges	UIR height	Feat. D.G. nodes	Feat. D.G. edges	Squashed D.G. edges	Atom D.T. nodes	Atom D.T. edges	Execution time	Memory usage
rustdesk	81 965	2826	108 908	8	1	67	8	1353	1346	13	24	41	27	29	22	1 s	53 MB
gitoxide	9490	76 429	223 296	79	2	554	139	180 166	180 089	53	321	910	473	587	510	817 s	999 MB
deno	101 658	419 085	284 483	36	3	561	8	123 883	123 850	31	124	1220	153	570	537	320 s	945 MB
tauri	89 324	3840 655	82 089	26	2	255	65	26 175	26 151	30	78	193	89	144	120	108 s	902 MB
sway	62 453	1092	210 721	28	4	369	7	119 724	119 700	33	75	118	78	91	67	1156 s	1910 MB
fuel-core	57 814	265 047	148 682	36	4	370	69	85 528	85 496	31	132	437	188	351	319	153 s	488 MB
alacrity	57 640	192 714	32 812	5	1	39	6	27 244	27 240	24	17	47	25	36	32	117 s	513 MB
zed	54 095	54 231	591 481	178	2	2458	91	23 791	23 615	28	381	268	237	231	55	91 s	497 MB
bat	51 048	1 496 904	14 971	1	0	30	9	675	674	19	5	14	7	8	7	26 s	339 MB
ripgrep	50 241	938 719	50 285	10	2	40	8	60 181	60 173	27	43	158	69	124	116	320 s	887 MB
meilisearch	49 152	1288	166 851	19	2	196	30	44 000	43 983	30	46	69	41	55	38	404 s	2477 MB
fuels-rs	43 907	5611	42 484	24	2	90	22	43 578	43 556	28	70	165	78	138	116	161 s	691 MB
typst	37 358	57 580	112 435	20	4	185	15	48 913	48 897	27	49	79	50	57	41	81 s	465 MB
helix	35 745	103 645	93 042	14	1	143	12	55 807	55 794	28	59	175	89	133	120	595 s	980 MB
ruff	35 669	4764	370 658	38	3	402	23	124 392	124 357	51	130	237	161	176	141	462 s	504 MB
lapce	34 961	35 019	67 802	5	2	73	3	20 525	20 522	41	16	44	25	32	29	195 s	948 MB
nushell	33 821	975	291 737	36	6	213	26	126 096	126 066	28	108	303	138	232	202	662 s	956 MB
polars	31 805	2400 845	376 212	26	3	362	789	58 368	58 345	32	103	448	161	354	331	277 s	529 MB
swc	31 649	1 767 375	648 247	116	15	957	230	412 313	412 212	258	326	1122	369	722	621	2016 s	3425 MB
influxdb	29 458	283 996	54 304	19	2	352	15	62 082	62 065	27	58	116	67	92	75	164 s	504 MB
tabby	29 742	5637	41 457	19	1	270	15	52 972	52 954	21	55	134	56	109	91	442 s	1827 MB
servo	29 222	8792	367 323	8	1	45	19	2670	2663	10	19	23	17	17	10	1 s	64 MB
wasmer	19 335	4820 432	263 877	35	8	350	153	72 860	72 833	26	121	318	164	235	208	246 s	435 MB
diem	16 702	16 759	428 331	186	10	2239	130	427 311	427 135	65	477	1207	475	757	581	1359 s	968 MB
texture-synthesis	1768	61 696	4735	3	0	7	2	13 211	13 208	23	15	30	22	21	18	11 s	295 MB
kajiya	5006	1060	26 638	14	3	97	3	28 922	28 911	29	25	22	17	19	8	68 s	1395 MB
rust-gpu	7412	2135	44 422	19	1	66	24	7581	7563	34	46	41	38	31	13	12 s	360 MB
substrate	8381	1938	595 987	270	8	2986	554	516 169	515 907	43	863	2172	1002	1681	1419	1135 s	1801 MB
tantivy	12 662	5253 776	118 896	9	1	67	11	34 915	34 907	51	32	84	41	66	58	104 s	478 MB
tonic	10 477	93 238 823	41 456	29	0	166	42	37 584	37 555	25	88	330	104	254	225	270 s	545 MB
sendme	357	15 473	1031	1	0	20	0	1577	1576	18	2	1	1	1	0	52 s	525 MB
komodo	2766	822	63 829	12	3	57	2	12 443	12 434	33	20	33	13	31	22	80 s	685 MB
quiche	9841	541 561	84 162	9	2	28	0	23 455	23 448	67	25	62	30	50	43	28 s	352 MB
rolldown	10 119	941	49 343	35	1	256	12	35 014	34 980	32	89	149	76	126	92	64 s	354 MB
iroh	3892	96 831	39 197	7	1	129	14	15 758	15 752	20	22	70	28	43	37	100 s	381 MB
sp1	1194	1238	104 219	25	1	363	44	70 776	70 752	53	77	183	87	147	123	248 s	549 MB
union	22 086	11 795	387 611	148	2	1562	248	134 318	134 172	33	518	843	605	605	459	212 s	563 MB
hyperswitch	13 381	13 561	575 328	33	6	450	116	107 187	107 160	29	123	754	183	578	551	599 s	3800 MB
egui	23 706	4248 075	101 376	39	4	138	44	39 574	39 539	44	111	235	126	181	146	88 s	489 MB
pueue	5297	72 209	18 340	3	2	24	0	15 463	15 462	23	5	18	8	15	14	64 s	506 MB
Total	1,212,599	120,362,360	7,329,058	1,628	116	17,036	3,008	3,294,554	3,293,042	258	4,898	12,873	5,618	9,129	7,617	13,328 s	3,800 MB
Average	30,314	3,009,059	183,226	40	2	425	75	82,363	82,326	37	122	321	140	228	190	333 s	885 MB




Figure 3. Correlation between lines of code (x-axis), execution time (s) ●, and peak memory usage (MB) ●.

capacity of mainstream laptops. At the crate level, average peak memory dropped to 105 MB, indicating that memory requirements depend more on the internal complexity of each crate than on the overall project size. This stability across a diverse set of projects confirms that RUSTYEx can be executed in environments with limited resources, such as cloud-based CI/CD pipelines or developer laptops.

Intermediate Structures. The intermediate representations produced by RUSTYEx remained compact and efficient to process. On average, the UIR contained approximately 82,000 nodes and a similar number of edges, which is relatively small compared to the size of the original ASTs. The *feature dependency graphs* were significantly more compact, typically comprising between 100 and 500 edges after multi-edge squashing. Even smaller were the atom dependency trees, which were on average 95% smaller than the UIR itself. This compactness is particularly important because it directly impacts the cost of subsequent analyses and ranking procedures: smaller graphs and trees allow for faster computation of centrality measures and logical transformations. These results highlight that our abstraction strategy, centered on atoms and UIR, strikes a good balance between preserving semantic detail and reducing analysis overhead.

Feature Complexity. Most projects defined a substantial number of features, with a median of 425 per project. On average, RUSTYEx detected 122 nodes in the feature dependency graph, often revealing cross-feature dependencies that underscore the challenges of analyzing highly configurable systems. As shown in Fig. 4, the number of detected features is generally greater than the number explicitly declared in Cargo.toml files. This discrepancy arises mainly from two factors: 1. feature overlaps across different crates within the same workspace, which are not redefined in the manifest, and 2. the handling of the not predicate, which introduces logically independent variants (e.g., for a feature *a*, the predicate `not(a)` is tracked as a separate node $\neg a$; in the worst case this doubles the detected feature count).

Despite the substantial number of features, their utilization in terms of *frequency* remained modest: on average, only 228 feature-related artifacts were identified, which means that each feature was used fewer than two times on average. We note, however, that a feature used infrequently may still govern a large block of code; this effect is precisely what the weight-based ADT refinement captures. This finding nonetheless motivates the need for principled prioritization: if many features are rarely exercised, developers and testers should concentrate on those that have the greatest structural influence.

Answer to RQ2. RUSTYEx scales to projects with hundreds of thousands of lines of code and thousands of features, with a median analysis time of 163 s per project and a peak memory footprint of 3.8 GB for the largest codebases.

4.4. RQ3: Effectiveness

Experimental Setup. We evaluate effectiveness by comparing RUSTYEx-prioritized configurations against two baselines: (i) *random sampling* and (ii) *feature-model-only centrality* (FDG-based ranking without ADT refinement, hereafter FM-only). The comparison metric is mutation kill rate, measured using cargo-mutants (Pool, 2022).




Figure 4. Delta between the number of *declared features* and the number of *detected features*.

Table 2

Cumulative mutation kill rate (%) at each configuration step k for the three strategies on bench-numeric (292 total mutants, 262 catchable).

k	RustyEx	FM-only	Random (expected)
1	20.5%	30.5%	17.9%
2	51.0%	46.6%	35.9%
3	66.4%	67.1%	53.8%
4	82.5%	82.5%	71.8%
5	89.7%	89.7%	89.7%

We designed a controlled benchmark crate, bench-numeric, included in the replication package,¹⁴ implementing five independent numerical modules each guarded by a dedicated Cargo feature: *statistics* (12 functions, 17 unit tests), *linear_algebra* (12 functions, 15 tests), *combinatorics* (8 functions, 16 tests), *geometry* (7 functions, 8 tests), and *trigonometry* (5 functions, 6 tests). With all features enabled, cargo-mutants generated 292 mutants (timeout 15 s per mutant, 8 parallel workers). For each strategy we generated a budget of $K = 5$ configurations using an incremental design: configuration k enables the top- k features according to the strategy's ranking. All three strategies are evaluated with the same cargo-mutants setup and random seed.

1. **RustyEx:** features ranked by Katz centrality with ADT refinement.
2. **FM-only:** Katz centrality on the FDG alone, without ADT weights (ablation via `-no-adt`).
3. **Random:** expected cumulative kill rate over all $K!$ orderings, computed as $\frac{k}{K} \cdot C_{\text{total}}$ where C_{total} is the total catchable-mutant count.

Results. Table 2 shows the cumulative mutation kill rate at each step $k = 1, \dots, 5$. Both centrality-based strategies consistently outperform the random baseline. RUSTYEx ranks *combinatorics* first (Katz = 0.305) because iterative and recursive algorithms (factorial, prime sieve, GCD) produce the highest ADT node count; FM-only ranks *statistics* first (Katz = 0.472) because mean is called by four sibling functions, making it the most central FDG node. At $k = 2$, RUSTYEx reaches 51.0% kill rate versus 46.6% for FM-only and 35.9% expected for Random. At $k = 5$ (full budget), all strategies converge to 89.7%, as all 262 catchable mutants are killed.

Answer to RQ3. Both RUSTYEx and FM-only outperform random sampling across all configuration budgets. At $k = 2$, RUSTYEx achieves a 15.1 percentage-point advantage over the random baseline. The ADT weight refinement correctly identifies the most structurally complex features; whether this translates to a consistent margin over FM-only depends on the project's testing strategy and is left for future evaluation on a larger project set.

¹⁴ Step-by-step reproduction instructions are in benchmarks/rq3_benchmarks/README.md; the single entry point is `python3 rq3-effectiveness.py` run from benchmarks/. Prerequisites: rustup with nightly-2025-02-20 and cargo-mutants ≥ 27 .

Table 3

Breakdown of the 116 crate failures by root cause.

Failure Category	Count	% of 116
Missing system dependencies	101	87.1%
Timeout (> 10 min)	5	4.3%
Other (parse error, ICE)	10	8.6%
Total	116	100%

4.5. RQ4: Robustness

Robustness. Out of more than 1600 crates analyzed (roughly 40 per project), only 116 failed, yielding a 93% success rate. Table 3 breaks down the 116 failures by root cause, as determined by inspecting the error output in the replication package. Most failures were attributable to missing system dependencies during compilation—a limitation of the host environment rather than the tool itself. Timeouts were rare, affecting fewer than 2% of all crates. Importantly, no major crashes or incorrect results were observed. Even large modular workspaces such as `diem` and `deno` were analyzed successfully without manual intervention.

These results confirm that RustyEx is mature and reliable enough to be integrated into the toolchains of both researchers and practitioners. **Answer to RQ4.** RustyEx achieves a 93% success rate across 1628 crates. The dominant failure mode is missing host-level system dependencies, which is independent of the tool's correctness; timeouts account for fewer than 2% of all analyzed crates.

5. Discussion

We discuss the key findings of our evaluation, the practical implications of RustyEx, and the limitations of the current work.

5.1. Practical adoption scenarios

RustyEx is designed for periodic, heavyweight analysis rather than every-commit feedback. Our results suggest three natural adoption scenarios:

1. **Nightly or pre-release analysis.** The median project analysis time of 163 s (≈ 2.7 min) makes RustyEx suitable for nightly CI jobs, where a longer window is typically available and thoroughness is more valuable than latency (Hilton et al., 2016). The 33-minute peak (for `swc` at 648 K LoC) is exceptional and is driven by the sheer scale of the project.
2. **Release-gate configuration auditing.** Before a major release, teams can run RustyEx once to identify the most relevant configurations and prioritize which ones to test or profile. The resulting ranked list remains valid as long as the feature dependency structure does not change substantially.
3. **Focused performance and security analysis.** As noted in Section 3.8, prioritized configurations can guide targeted profiling or vulnerability analysis on the configurations most likely to exercise critical code paths.

For every-commit feedback, RustyEx is not the right tool; standard test runners and lightweight linters serve that purpose better.

5.2. Implications for researchers

Our results reveal several non-obvious aspects of real-world Rust variability.

Feature utilization is sparse. Despite a median of 425 declared features per project, the average number of feature-related artifacts is only 228. This means that the configuration space is vast but sparsely exercised: most features influence little code in isolation. However,

their interactions—captured by the feature dependency graph—can be non-trivial. This motivates centrality-based prioritization: degree-based or random selection would assign equal probability to rarely-exercised features, diluting the testing effort.

Compiler-level analysis reaches features invisible to feature models. Because RustyEx operates on the compiler's AST, it detects `cfg` attributes that a feature model or manifest-level analysis would miss, such as platform-specific guards (`#[cfg(unix)]`) that are not declared as named Cargo features. These contribute to the detected-vs.-declared feature discrepancy observed in Fig. 4 and represent a genuine coverage advantage over FM-only approaches (Hentze et al., 2022; Krieter et al., 2022).

ADT refinement vs. FDG-only ranking. A natural question is whether the atom dependency tree refinement materially improves over FDG-only centrality. This is precisely the question addressed by RQ3 (Section 4.4). Pending those results, we note theoretically that the refinement matters most for projects where features have highly unequal code extents—i.e., where some features gate large modules while others gate single statements. In our dataset, this is the common case (the ratio of UIR nodes to feature dependency graph nodes averages 675:1), suggesting that weight-based refinement is likely to reorder the ranking non-trivially.

5.3. Limitations and open challenges

No ground truth for configuration importance. There is no universally accepted benchmark for which configurations are “most important”. RQ3 uses mutation kill rate as a proxy for fault-detection effectiveness, but this metric has its own limitations: surviving mutants may be equivalent, and the set of mutations generated by cargo-mutants may not be representative of real faults. Establishing a richer benchmark suite for Rust configuration testing remains an open challenge.

Language scope. This work focuses on Rust, whose `cfg` system differs structurally from C/C++ preprocessor macros and Java conditional annotations. Extending RustyEx to other languages requires reimplementing the compiler plugin layer; the formalization of UIR, FDG, and ADT is language-agnostic.

Mutual exclusion and feature interactions. The current CNF formulation encodes feature dependencies but does not model mutual exclusion constraints beyond those derivable from `Cargo.toml`. In projects where features are implicitly mutually exclusive at the code level (e.g., two platform-specific implementations of the same interface), the SAT solver may generate configurations that compile successfully but lead to undefined behavior at runtime. Detecting and encoding such constraints automatically is left for future work.

Dynamic variability. RustyEx analyzes compile-time variability only. Runtime configuration via environment variables, feature flags, or plugin architectures is outside the current scope.

6. Threats to validity

We organize our discussion following Wohlin et al. (2012)'s taxonomy.

6.1. Construct validity

Proxy metrics for feature relevance. We rely on *geometric* and *spectral* centrality computed on the feature dependency graph to approximate feature relevance. These metrics are not neutral: each has biases and limitations, especially regarding disconnected components and local vs. global influence. In particular, eigenvector centrality is undefined on disconnected graphs and closeness centrality becomes degenerate in the presence of isolated nodes.

Mitigation. In our evaluation we addressed this threat in two concrete ways. First, we introduced a *patch node* that connects all otherwise

isolated components, ensuring every feature participates in the centrality computation. Second, we ran the evaluation with all four supported centrality measures (Newman closeness, Opsahl betweenness, eigenvector, and Katz) and verified that the relative ordering of the top-ranked features was consistent across measures for the majority of projects. Users can reproduce this sensitivity analysis using the `-centrality` flag of `RUSTyEx`.

Weighting of UIR nodes. In cases where definitions are missing or cycles are detected, nodes of type *reference* are assigned a default fallback weight. This approximation may underestimate the role of unresolved external calls, potentially distorting prioritization.

Mitigation. In our evaluation, cycles arose primarily from direct and mutual recursion, which `RUSTyEx` detects via fixed-point termination. We measured that fewer than 1% of UIR nodes in the dataset received a fallback weight, and we verified manually on three representative projects that the fallback assignment did not alter the top-5 ranked features. The replication package includes the list of affected nodes per project for independent verification.

6.2. Internal validity

Assumptions about SAT solver correctness. Our configuration generation relies on the soundness and completeness of the underlying SAT solver. If the solver produced incorrect results, the validity of generated configurations would be at risk.

Mitigation. We used three well-tested, widely deployed SAT solvers (MiniSat, CaDiCaL, Kissat) and cross-validated their outputs on a sample of 10 projects: in all cases the configurations produced by different solvers were logically equivalent (i.e., they satisfied the same set of clauses). Additionally, the formal correctness argument in Section 3.6 holds independently of the specific solver implementation.

Timeouts and analysis failures. A fraction of crates (~ 7%) failed due to timeouts or missing system dependencies. Such failures may introduce selection bias, as not all crates are equally represented in the analysis.

Mitigation. We applied a uniform 10-minute timeout per crate across all projects, ensuring no project received preferential treatment. We further verified that the projects with the highest failure counts (e.g., `nushell`: 6 failures, `swc`: 15 failures) still had the majority of their crates analyzed successfully, and that failed crates belonged to peripheral sub-packages rather than core modules. The robustness breakdown in Table 3 (Section 4.5) shows that most failures stem from missing host-level system dependencies rather than from intrinsic tool limitations.

6.3. External validity

Focus on open-source Rust projects. Our evaluation is based on 40 open-source projects from GitHub and `crates.io`. While these projects cover a broad spectrum of real-world software domains (compilers, databases, network stacks, UI frameworks, cryptographic libraries), they may not capture the full variability of proprietary or industrial codebases.

Mitigation. We selected projects by GitHub star count and `crates.io` download rank to maximize representativeness across application domains and codebase sizes (from 1 K to 650 K LoC). Many of the analyzed projects are widely used in production (e.g., `deno`, `ripgrep`, `meilisearch`) and serve as dependencies for industrial systems, increasing confidence in external generalizability. We also note that the approach is language-agnostic; future work will evaluate it on C/C++ and other ecosystems with native variability support.

Fixed configuration budget. We generated a fixed number K of configurations per project. In practice, real-world scenarios may require adaptive strategies that vary the number of generated configurations depending on project size, release stage, or available resources.

Mitigation. We selected K conservatively (set to 10) to ensure the generation step completed within the per-crate timeout on all analyzed projects. In the replication package, we provide scripts to re-run the analysis with different values of K , allowing practitioners to explore the trade-off between configuration budget and coverage.

6.4. Conclusion validity

Dependence on feature ranking. The order in which configurations are generated depends on both the selected centrality metric and the refinements applied by the atom dependency tree. Different choices may therefore lead to different prioritized sets.

Mitigation. In our evaluation we ran all four centrality measures and compared the resulting rankings. We found that the top-3 ranked features agreed across all three measures in 100% of our validation test cases, suggesting the ranking is robust to metric choice for the most central features. The sensitivity data is included in the replication package.

No ground truth for configurations. There is no universally accepted reference set of “correct” or “important” configurations. As a result, it is challenging to directly evaluate the relevance of the configurations produced by our method.

Mitigation. We addressed this by pursuing two complementary forms of validation: (i) structural soundness—all generated configurations satisfy the CNF formula derived from `Cargo.toml` constraints, verified exhaustively across 1628 crates (RQ1); and (ii) empirical effectiveness—RQ3 (Section 4.4) compares configurations produced by `RUSTyEx` against random and FM-only baselines using mutation kill rate as an objective proxy for fault-detection capability.

7. Related work

Prioritizing configurations in highly configurable systems has been extensively studied, though from different perspectives and often under different assumptions. Several comprehensive surveys provide an overview of the field, such as Agh et al. (2024) (see in particular Sect. 4.4), El-Sharkawy et al. (2019), and Idham et al. (2023), which systematically map research trends and highlight open challenges. Building on these, we briefly summarize the most relevant contributions along four main dimensions.

Static Analysis and Preprocessing. Early work on highly configurable systems, especially in the context of the Linux kernel, focused on static analysis and preprocessing techniques. El-Sharkawy et al. (2020) developed methods to process `#ifdef` directives for extracting software metrics, making configuration-specific complexity more manageable. Similarly, Sincero et al. (2010) investigated preprocessing of C macros to detect dead code, thus helping to reduce variability-induced maintenance effort. The TypeChef framework (Kenner et al., 2010) pioneered variability-aware type checking by lifting the C type system to work over presence conditions derived from `#ifdef` directives, enabling analysis without fixing a configuration. The KernelHaven workbench (Kröher et al., 2018) later generalized this idea into a plugable platform for extracting and analyzing feature constraints from large C-based product lines. These approaches are closely related to our idea of analyzing variability from the source, but they remain tied to metrics extraction, defect detection, or specific ecosystems. By contrast, `RUSTyEx` generalizes variability-aware static analysis beyond macros or directives, targeting the identification of *relevant configurations* in Rust projects without assuming language-specific preprocessing artifacts.

Solution-Space and Presence-Condition Sampling. A growing body of work moves beyond feature models to incorporate source-level information in sampling. Hentze et al. (2022) proposed solution-space sampling for multi-domain product lines, generating configurations by reasoning over the set of valid product variants without requiring an explicit feature model. Krieter et al. (2022) developed T-wise presence-condition coverage, extending t-wise interaction testing to

presence conditions extracted directly from source artifacts rather than from feature diagrams. These works demonstrate that considering the solution space improves sampling quality compared to feature-model-only approaches. *RUSTyEx* differs in a complementary but distinct way: rather than sampling to achieve combinatorial coverage of presence conditions, it *rank*s configurations by their structural importance, using centrality measures on a code-weighted dependency graph. The two strategies are orthogonal and could be combined: *RUSTyEx*'s ranking could seed the initial configurations for a subsequent T-wise coverage pass.

Feature Model-based Prioritization. Another major line of research builds on explicit feature models, where features and their relationships are captured in structured diagrams. Bagheri et al. (2010) proposed the *stratified analytic hierarchy process* to prioritize and select features by decomposing the decision process into manageable layers. Peng et al. (2015) employed a directed, weighted acyclic graph to assess feature importance via *weighted degree centrality*, highlighting asymmetries in feature influence. Mannion and Kaindl (2021) explored weighting strategies based on variability types, assuming a graph with explicit sources and sinks. Beyond prioritization, Bagheri and Gasevic (2011) also assessed the maintainability of feature models using structural metrics such as cyclomatic complexity (McCabe, 1976), network centrality (Newman, 2010), and classical software engineering metrics (Fenton, 1991). While these approaches are rigorous and effective, they presuppose the existence of a formal feature model—a strong assumption in practice, since many modern systems (including Rust projects) rely on decentralized, implicit forms of variability encoded directly in build manifests and conditional compilation. *RUSTyEx* differs in that it does not require a predefined feature model, instead reconstructing variability information directly from source and build metadata.

Centrality Measures for Prioritization. A complementary strand of research applies graph-theoretic concepts, especially centrality measures, to variability management and testing. Mohammed et al. (2024) ranked configurations within a single feature model using a variety of centrality metrics, thus identifying configurations with disproportionately high influence. Levasseur and Badri (2024) used centrality, object-oriented metrics, and machine learning to prioritize unit tests, showing that structural graph properties can guide resource allocation in testing. Ferreira et al. (2016) went further by proposing *variational call graphs* (Ferreira et al., 2015; Johnson, 2009) enriched with centrality information to pinpoint functions that may become vulnerabilities under certain feature combinations. While all these works demonstrate the value of centrality in highlighting critical elements, they remain tied either to explicit feature models, to testing strategies, or to specialized tasks such as vulnerability detection. Our approach builds on the same intuition—that centrality can capture relevance—but applies it to a new abstraction level: the atom dependency tree derived from Rust's variability constructs, enriched with code-extent weights that none of the above works consider. In doing so, *RUSTyEx* extends the applicability of centrality-based prioritization to ecosystems where feature models are implicit and configuration spaces are both large and sparsely exercised.

8. Conclusion

In this paper, we presented the first general method for prioritizing configurations in highly configurable software systems via a compiler-based refined ranking of features. Unlike previous approaches, which are predominantly based on explicit feature models or consider source artifacts without code-extent weighting (Hentze et al., 2022; Krieter et al., 2022), our method does not require a predefined feature model and explicitly accounts for both feature dependencies and the extent of code affected by each feature.

The method combines inter-procedural static analysis to extract a *unified intermediate representation* (UIR), automated construction of the

feature dependency graph (FDG) and the *atom dependency tree* (ADT), centrality-based ranking with ADT-driven weight refinement, and CNF-based SAT solving for configuration generation. The soundness of the approach is formally guaranteed by construction (Section 3.6) and confirmed empirically across 1628 crates with a 100% validity rate.

We implemented the method in *RUSTyEx*, a fully automated tool for Rust software. Our evaluation (Section 4) demonstrates that *RUSTyEx* is scalable—handling projects up to 650 K LoC with a median analysis time of 163 s—and robust, achieving a 93% success rate across a diverse set of 40 open-source projects. The approach is language-agnostic and can in principle be applied to other ecosystems with native variability support, such as C/C++ and Java.

Open Challenges. Two key directions remain open for future work. First, the effectiveness of the centrality-guided ranking (RQ3, Section 4.4) requires a rigorous empirical study comparing *RUSTyEx*-prioritized configurations against random and FM-only baselines in terms of fault-detection capability—an experiment currently in progress. Second, extending *RUSTyEx* to other language ecosystems requires reimplementing the compiler plugin layer; the formalization of UIR, FDG, and ADT is language-agnostic, and future work will investigate applications to C/C++ product lines using infrastructure such as KernelHaven (Kröher et al., 2018) or TypeChef (Kenner et al., 2010). Finally, incorporating solution-space-aware sampling strategies (Hentze et al., 2022; Krieter et al., 2022) alongside centrality-guided prioritization represents a promising research direction for achieving both structural coverage and feature relevance in a unified framework.

CRedit authorship contribution statement

Federico Bruzzone: Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Walter Cazzola:** Writing – review & editing, Writing – original draft, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Luca Favini:** Writing – review & editing, Writing – original draft, Software, Methodology, Conceptualization.

Declaration of competing interest

We wish to confirm that there are not known conflicts of interest associated with this publication and there has been no significant financial support for this work that could have influenced its outcome.

Data availability

A replication package is available on zenodo: <https://zenodo.org/records/21172650>.

References

- Agh, H., Azamnour, A., Wagner, S., 2024. Software product line testing: A systematic literature review. *Empir. Softw. Eng.* 29.
- Al-Hajjaji, M., Thüm, T., Lochau, M., Meinicke, J., Saake, G., 2019. Effective product-line testing using similarity-based product prioritization. *Softw. Syst. Model.* 18, 499–521.
- Anderson, B., Bergstrom, L., Goregaokar, M., Matthews, J., McAllister, K., Moffitt, J., Sapin, S., 2016. Engineering the servo web browser engine using rust. In: Xie, T., Zhang, D. (Eds.), *Proceedings of the 38th International Conference on Software Engineering: Software Engineering in Practice. ICSE-SEIP'16*, IEEE, Austin, TX, USA, pp. 81–89.
- Apel, S., Batory, D., Kästner, C., Saake, G., 2013. *Feature-Oriented Software Product Lines*. Springer.
- Apel, S., Speidel, H., Wendler, P., von Rhein, A., Beyer, D., 2011. Detection of feature interactions using feature-aware verification. In: Păsăreanu, C., Hosking, J. (Eds.), *Proceedings of the 26th International Conference on Automated Software Engineering. ASE'11*, IEEE, Lawrence, KS, USA, pp. 372–375.
- Bagheri, E., Asadi, M., Gasevic, D., Soltani, S., 2010. Stratified analytic hierarchy process: Prioritization and selection of software features. In: Bosch, J., Lee, J. (Eds.), *Proceedings of the 14th International Software Product Line Conference. SPLC'10*, Springer, Jeju Island, South Korea, pp. 300–315.

- Bagheri, E., Gasevic, D., 2011. Assessing the maintainability of software product line feature models using structural metrics. *Softw. Qual. J.* 19, 576–612.
- Bavelas, A., 1950. Communication patterns in task-oriented groups. *J. Acoust. Soc. Am.* 22, 725–730.
- Benavides, D., Segura, S., Ruiz-Cortés, A., 2010. Automated analysis of feature models 20 years later: A literature review. *Inf. Syst.* 35, 615–636.
- Bergstra, J.A., Tucker, J.V., 1995. Equational specifications, complete term rewriting systems, and computable and semicomputable algebras. *J. ACM* 42, 1194–1230.
- Berman, A., Plemmons, R.J., 1979. Nonnegative Matrices in the Mathematical Sciences. Academic Press.
- Biere, A., Fazekas, K., Fleury, M., Heisinger, M., 2020. CaDiCaL, kissat, paracooba, pilinging and treengeling entering the SAT competition 2020. In: Balyoi, T., Froyleyks, N., Heule, M.J.H., Iser, M., Järvisalo, M., Suda, M. (Eds.), *Proceedings of the SAT Competition 2020*. University of Helsinki, Alghero, Italy, pp. 50–53.
- Bloch, F., Jackson, M.O., Tebaldi, P., 2023. Centrality measures in networks. *Soc. Choice Welf.* 61, 413–453.
- Boldi, P., Vigna, S., 2014. Axioms for centrality. *Internet Math.* 10, 222–262.
- Bonacich, P., 1972. Factoring and weighting approaches to status scores and clique identification. *J. Math. Sociol.* 2, 113–120.
- Borgatti, S.P., 2005. Centrality and network flow. *Soc. Netw.* 25, 55–71.
- Chen, S.F., Wu, Y.S., 2022. Linux kernel module development with rust. In: Aafer, Y., Li, S., Wu, Y. (Eds.), *Proceedings of the Conference on Dependable and Secure Computing, DSC'22*, IEEE, Edinburgh, United Kingdom, pp. 1–2.
- Classen, A., Cordy, M., Schobbens, P.Y., Heymans, P., Legay, A., Raskin, J.F., 2013. Featured transition systems: Foundations for verifying variability-intensive systems and their application to LTL model checking. *IEEE Trans. Softw. Eng.* 39, 1069–1089.
- Clements, P., Northrop, L., 2001. *Software Product Lines: Practices and Patterns*. Addison-Wesley.
- Cohen, M.B., Dwyer, M.B., Shi, J., 2007. Interaction testing of highly-configurable systems in the presence of constraints. In: Rosenblum, D.S., Elbaum, S.G. (Eds.), *Proceedings of the 16th International Symposium on Software Testing and Analysis, ISSTA'07*, ACM, London, United Kingdom, pp. 129–139.
- Cooper, K.D., Torczon, L., 2022. *Engineering a Compiler*. Morgan Kaufmann.
- Das, K., Samanta, S., Pal, M., 2018. Study on centrality measures in social networks: A survey. *Soc. Netw. Anal. Min.* 8, 13:1–13:11.
- Dasgupta, S., 1999. Learning polytrees. In: Laskey, K.B., Prade, H. (Eds.), *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence, UAI'99*, Morgan Kaufmann Publishers Inc., Stockholm, Sweden, pp. 134–141.
- Eén, N., Sörensson, N., 2003. An extensible SAT-solver. In: Giunchiglia, E., Tacchella, A. (Eds.), *Proceedings of the 6th International Conference on Theory and Applications of Satisfiability Testing, SAT'03*, Springer, Santa Margherita Ligure, Italy, pp. 502–518.
- El-Sharkawy, S., Krafczyk, A., Schmid, K., 2017. An empirical study of configuration mismatches in linux. In: Cohen, M., Acher, M. (Eds.), *Proceedings of the 21st International Systems and Software Product Line Conference, SPLC'17*, ACM, Sevilla, Spain, pp. 19–28.
- El-Sharkawy, S., Krafczyk, A., Schmid, K., 2020. Fast static analyses of software product lines: An example with more than 42,000 metrics. In: Cordy, M., Acher, M. (Eds.), *Proceedings of the 14th International Working Conference on Variability Modelling of Software-Intensive Systems (VaMoS'20)*, ACM, Magdeburg Germany, pp. 1–9.
- El-Sharkawy, S., Yamagishi-Eichler, N., Schmid, K., 2019. Metrics for analyzing variability and its implementation in software product lines: A systematic literature review. *Inf. Softw. Technol.* 106, 1–30.
- Fenton, N.E., 1991. *Software Metrics: a Rigorous Approach*. Chapman & Hall, London.
- Ferreira, G., Kästner, C., Pfeffer, J., Apel, S., 2015. Characterizing complexity of highly-configurable systems with variational call graphs: Analyzing configuration options interactions complexity in function calls. In: *Proceedings of the Symposium and Bootcamp on the Science of Security, HotSoS'15*, ACM, Urbana-Champaign, IL, USA, pp. 1–2.
- Ferreira, G., Malik, M., Kästner, C., Pfeffer, J., Apel, S., 2016. Do #ifdefs influence the occurrence of vulnerabilities? An empirical study of the linux kernel. In: Rabiser, R., Xie, B. (Eds.), *Proceedings of the 20th International Systems and Software Product-Line Conference, SPLC'16*, ACM, Beijing, China, pp. 65–73.
- Freeman, L.C., 1977. A set of measures of centrality based on betweenness. *Sociometry* 40, 35–41.
- Gamma, E., Helm, R., Johnson, R., Vlissides, J., 1995. *Design Patterns: Elements of Reusable Object-Oriented Software*. Professional Computing Series. Addison-Wesley, Reading, MA, USA.
- Halin, A., Nuttinck, A., Acher, M., Devroey, X., Perrouin, G., Baudry, B., 2019. Test them all, is it worth it? Assessing configuration sampling on the JHipster web development stack. *Empir. Softw. Eng.* 24, 674–717.
- Hentze, M., Pett, T., Sundermann, C., Krieter, S., Thüm, T., Schaefer, I., 2022. Generic solution-space sampling for multi-domain product lines. In: Kameyama, Y. (Ed.), *Proceedings of the 21st International Conference on Generative Programming: Concepts and Experiences, GPCE'22*, ACM, Auckland, New Zealand, pp. 135–147.
- Hilton, M., Tunnell, T., Huang, K., Marinov, D., Dig, D., 2016. Usage, costs, and benefits of continuous integration in open-source projects. In: Apel, S., Khurshid, S. (Eds.), *Proceedings of the 31st International Conference on Automated Software Engineering, ASE'16*, ACM, Singapore, Singapore, pp. 426–437.
- Howson, C., 1997. *Logic with Trees: An Introduction to Symbolic Logic*. Routledge.
- Idham, M., Abd Halim, S., Jawawi, D.N.A., Zakaria, Z., Erianda, A., Arss, N., 2023. Test case prioritization for software product line: A systematic mapping study. *J. Inform. Vis.* 7, 2126–2134.
- Johnson, N.F., 2009. *Simply Complexity: A Clear Guide to Complexity Theory*. Newworld Publications.
- Kang, K.C., Cohen, S.G., Hess, J.A., Novak, W.E., Peterson, A.S., 1990. *Feature-Oriented Domain Analysis (FODA) Feasibility Study*. Technical Report CMU/SEI-90-TR-21, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA.
- Karataş, A.S., Oğuztüzün, H., Doğru, A., 2010. Global constraints on feature models. In: Cohen, D. (Ed.), *Proceedings of the 16th International Conference on Principles and Practice of Constraint Programming, CP'10*, Springer, St. Andrews, Scotland, pp. 537–551.
- Katz, L., 1953. A new status index derived from sociometric analysis. *Psychometrika* 18, 39–43.
- Kennedy, A., Russo, C.V., 2005. Generalized algebraic data types and object-oriented programming. In: Gabriel, R.P. (Ed.), *Proceedings of 19th ACM International Conference on Object-Oriented Programming Systems, Languages and Applications, OOPSLA'05*, ACM, San Diego, CA, USA, pp. 21–40.
- Kenner, A., Kästner, C., Haase, S., Leich, T., 2010. TypeChef: Toward type checking #ifdef variability in c. In: Apel, S., Batory, D., Czarnecki, K., Heidenreich, F., Kästner, C., Nierstrasz, O. (Eds.), *Proceedings of the 2nd International Workshop on Feature-Oriented Software Development, FOSD'10*, ACM, Eindhoven, The Netherlands, pp. 25–32.
- Krieter, S., Thüm, T., Schulze, S., Ruland, S., Lochau, M., Saake, G., Leich, T., 2022. T-Wise Presence Condition Coverage and Sampling for Configurable Systems. Technical Report arXiv:2205.15180, <http://dx.doi.org/10.48550/arXiv.2205.15180>.
- Kröher, C., El-Sharkawy, S., Schmid, K., 2018. KernelHaven: An experimentation workbench for analyzing software product lines. In: Chechik, M., Harman, M. (Eds.), *Proceedings of the 40th International Conference on Software Engineering, ICSE'18*, ACM, Gothenburg, Sweden, pp. 73–76.
- Krueger, C.W., 2006. New methods in software product line practice. *Commun. ACM* 49, 37–40.
- Landherr, A., Friedl, B., Heidemann, J., 2010. A critical review of centrality measures in social networks. *Bus. Inf. Syst. Eng.* 2, 371–385.
- Lee, J., Hwang, S., 2019. Combinatorial test design using design-time decisions for variability. *J. Softw. Eng. Knowl. Eng.* 29, 1141–1158.
- Lehmann, D.J., Smyth, M.B., 1981. Algebraic specification of data types: A synthetic approach. *J. Math. Syst. Theory* 14, 97–139.
- Levasseur, M.A., Badri, M., 2024. Prioritizing unit tests using object-oriented metrics, centrality measures, and machine learning algorithms. *Innov. Syst. Softw. Eng.* 1–27.
- Li, H., Guo, L., Yang, Y., Wang, S., Xu, M., 2024. An empirical study of rust-for-linux: The success, dissatisfaction, and compromise. In: Bagchi, S., Zhang, Y. (Eds.), *Proceedings of the USENIX Annual Technical Conference, USENIX'24*, Curran Associates, Inc., Santa Clara, CA, USA, pp. 425–443.
- Li, Z., Wang, J., Sun, M., Lui, J.C.S., 2021. MirChecker: Detecting bugs in rust programs via static analysis. In: Vigna, G., Shi, E. (Eds.), *Proceedings of the 2021 Conference on Computer and Communications Security, CCS'21*, ACM, Virtual, South Korea, pp. 2183–2196.
- Liebig, J., Apel, S., Janker, A., Garbe, F., Oster, S., 2017. Handling static configurability in refactoring engines. *Computer* 50, 44–53.
- Liebig, J., Janker, A., Garbe, F., Apel, S., Lengauer, C., 2015. Morpheus: Variability-aware refactoring in the wild. In: Canfora, G., Elbaum, S. (Eds.), *Proceedings of the 37th International Conference on Software Engineering, ICSE'15*, IEEE, Florence, Italy, pp. 380–391.
- Lochau, M., Oster, S., Goltz, U., Schürr, A., 2012. Model-based pairwise testing for feature interaction coverage in software product line engineering. *Softw. Qual. J.* 20, 567–604.
- Magnusson, E., Hedin, G., 2007. Circular reference attributed grammars—Their evaluation and applications. *Sci. Comput. Program.* 68, 21–37.
- Mannion, M., Kaindl, H., 2021. Using binary strings for comparing products from software-intensive systems product lines. In: Schaefer, I., ter Beek, M.H. (Eds.), *Proceedings of the 25th International Software Product Line Conference, SPLC'21*, ACM, Leicester, United Kingdom, pp. 257–266.
- Marchiori, M., Latora, V., 2000. Harmony in the small-world. *Phys. A* 285, 539–546.
- Martin, H., Acher, M., Alves Pereira, J., Jézéquel, J.M., 2021. A comparison of performance specialization learning for configurable systems. In: Schaefer, I., ter Beek, M.H. (Eds.), *Proceedings of the 25th International Software Product Line Conference, SPLC'21*, ACM, Leicester, United Kingdom, pp. 46–57.
- Matsakis, N.D., Klock, F.S., 2014. The rust language. *ACM SIGAda Lett.* 34, 103–104.
- Mazurkiewicz, A., 1995. Introduction to trace theory. In: Diekert, V., Rozenberg, G. (Eds.), *The Book of Traces*. World Scientific Publishing, pp. 3–41 (Chapter 1).
- McCabe, T.J., 1976. A complexity measure. *IEEE Trans. Softw. Eng.* 2, 308–320.
- Melo, J., Flesborg, E., Brabrand, C., Wasowski, A., 2016. A quantitative analysis of variability warnings in linux. In: Schaefer, I., Alves, V. (Eds.), *Proceedings of the 10th International Workshop on Variability Modelling of Software-Intensive Systems, VaMoS'10*, ACM, Salvador, Brazil, pp. 3–8.
- Mohammed, F., Mannion, M., Kaindl, H., Patterson, J., 2024. Evaluating the relative importance of product line features using centrality metrics. In: Mecella, M., Rensink, A. (Eds.), *Proceedings of the 19th International Conference on Software Technologies, ICSoft 2024*, SciTe Press, Dijon, France, pp. 469–476.


- Newman, M.E.J., 2001. Scientific collaboration networks. II. Shortest paths, weighted networks, and centrality. *Phys. Rev. E* 64, 016132.
- Newman, M.E.J., 2010. *Networks: An Introduction*, first ed. Oxford University Press.
- von Oheimb, D., 1999. Hoare logic for mutual recursion and local variables. In: Rangan, C.P., Raman, V., Ramanujam, R. (Eds.), *Proceedings of the 19th Conference on Foundations of Software Technology and Theoretical Computer Science. FSTTCS'99*, Springer, Chennai, India, pp. 168–180.
- Olender, K.M., Osterweil, L.J., 1990. Cecil: A sequencing constraint language for automatic static analysis generation. *IEEE Trans. Softw. Eng.* 16, 268–280.
- Olender, K.M., Osterweil, L.J., 1992. Interprocedural static analysis of sequencing constraints. *Trans. Softw. Eng. Methodol.* 1, 21–52.
- Opsahl, T., Agneessens, F., Skvoretz, J., 2010. Node centrality in weighted networks: Generalizing degree and shortest paths. *Soc. Netw.* 32, 245–251.
- Oster, S., Markert, F., Ritter, P., 2010. Automated incremental pairwise testing of software product lines. In: Bosch, J., Lee, J. (Eds.), *Proceedings of the 14th International Software Product Line Conference. SPLC'10*, Springer, Jeju Island, South Korea, pp. 196–210.
- Parejo, J.A., Sánchez, A.B., Segura, S., Ruiz-Cortés, A., Lopez-Herrejon, R.E., Egyed, A., 2016. Multi-objective test case prioritization in highly configurable systems: A case study. *J. Syst. Softw.* 122, 287–310.
- Patel, S., Gupta, P., Shah, V., 2013. Combinatorial interaction testing with multi-perspective feature models. In: Liu, Y., Pang (Eds.), *Proceedings of the 6th International Conference on Software Testing, Verification and Validation Workshops. ICSTW'13*, IEEE, Luxembourg City, Luxembourg, pp. 321–330.
- Peng, Z., Wang, J., He, K., Li, H., 2015. An approach for prioritizing software features based on node centrality in probability network. In: Kapitsaki, G., Almeida, E. (Eds.), *Proceedings of the 15th International Conference on Software Reuse. ICSR'15*, Springer, Limassol, Cyprus, pp. 106–121.
- Perrouin, G., Sen, S., Klein, J., Baudry, B., le Traon, Y., 2010. Automated and scalable T-wise test CASE generation strategies for software product lines. In: Cavalli, A.R., Ghosh, S. (Eds.), *Proceedings of the 3rd International Conference on Software Testing, Verification and Validation. ICST'10*, IEEE, Paris, France, pp. 459–468.
- Pierce, B.C., 2002. *Types and Programming Languages*. MIT Press.
- Pohl, R., Lauenroth, K., Pohl, K., 2011. A performance comparison of contemporary algorithmic approaches for automated analysis operations on feature models. In: pásáreanu, C., Hosking, J. (Eds.), *Proceedings of the 26th International Conference on Automated Software Engineering. ASE'11*, IEEE, Lawrence, KS, USA, pp. 313–322.
- Pool, M., 2022. Cargo-mutants: Mutation testing tool for Rust. Accessed: 2025. <https://github.com/sourcefrog/cargo-mutants>.
- Prehofer, C., 1997. Feature-oriented programming: A fresh look at objects. In: Akšit, M., Matsuoka, S. (Eds.), *Proceedings of the 11th European Conference on Object-Oriented Programming. ECOOP'97*, Springer, Helsinki, Finland, pp. 419–443.
- Qu, X., Cohen, M.B., Rothermel, G., 2008. Configuration-aware regression testing: An empirical study of sampling and prioritization. In: Zeller, A. (Ed.), *Proceedings of the 17th International Symposium on Software Testing and Analysis. ISSTA'08*, ACM, Seattle, WA, USA, pp. 75–86.
- Sánchez, A.B., Segura, S., Parejo, J.A., Ruiz-Cortés, A., 2017. Variability testing in the wild: The drupal CASE study. *Softw. Syst. Model.* 16, 173–194.
- Sánchez, A.B., Segura, S., Ruiz-Cortés, A., 2014. A comparison of test case prioritization criteria for software product lines. In: Williams, L., Wohlin, C. (Eds.), *Proceedings of 7th International Conference on Software Testing, Verification and Validation. ICST'14*, IEEE, Cleveland, OH, USA, pp. 41–50.
- Seeley, J.R., 1949. The net of reciprocal influence; a problem in treating sociometric data. *Can. J. Psychol.* 3, 234–240.
- Seidl, C., Winkelmann, T., Schaefer, I., 2016. A software product line of feature modeling notations and cross-tree constraint languages. In: Oberweis, A., Reussner, R.H. (Eds.), *Proceedings of the 13th Edition of Modellierung. Modellierung'16*, Springer, Karlsruhe, Germany, pp. 157–172.
- Siegmund, N., Grebhanh, A., Apel, S., Kästner, C., 2015. Performance-influence models for highly configurable systems. In: Harman, M., Heymans, P. (Eds.), *Proceedings of the 10th Joint Meeting on Foundations of Software Engineering. ESEC/FSE'15*, ACM, Bergamo, Italy, pp. 284–294.
- Sincero, J., Schirmeier, H., Schröder-Preikschat, W., Spinczyk, O., 2007. Is the linux kernel a software product line? In: van der Linden, F., Linnell, B. (Eds.), *Proceedings of the 2nd International Workshop on Open Source Software and Product Lines. SPLC-OSSPL'07*, Kyoto, Japan.
- Sincero, J., Tartler, R., Lohmann, D., Schröder-Preikschat, W., 2010. Efficient extraction and analysis of preprocessor-based variability. In: Järvi, J. (Ed.), *Proceedings of the 9th International Conference on Generative Programming and Component Engineering. GPCE'10*, ACM, Eindhoven, The Netherlands, pp. 33–42.
- Sörensson, N., Eén, N., 2005. MiniSat v1.13 – A SAT Solver with Conflict-Clause Minimization. In: Bacchus, F. (Ed.), *Proceedings of the 8th International Conference on Theory and Applications of Satisfiability Testing. SAT'05*, Springer, St. Andrews, Scotland, pp. 1–2, Poster.
- Souto, S., d'Amorim, M., 2018. Time-space efficient regression testing for configurable systems. *J. Syst. Softw.* 137, 733–746.
- Tartler, R., Dietrich, C., Sincero, J., Schröder-Preikschat, W., Lohmann, D., 2014. Static analysis of variability in system software: The 90,000 #ifdefs issue. In: Gibson, G., Zeldovich, N. (Eds.), *Proceedings of the USENIX Annual Technical Conference. USENIX ATC'14*, USENIX Association, Philadelphia, PA, USA, pp. 421–432.
- Tartler, R., Lohmann, D., Sincero, J., Schröder-Preikschat, W., 2011. Feature consistency in compile-time-configurable system software: Facing the linux 10,000 feature problem. In: Heiser, G. (Ed.), *Proceedings of the 6th Conference on Computer Systems. EuroSys'11*, ACM, Salzburg, Austria, pp. 47–60.
- The Linux Kernel Organization, 2021. Rust in the linux kernel. Accessed: 2025. <https://docs.kernel.org/rust/index.html>.
- Treinish, M., Carvalho, I., Tsimikounakis, G., Sá, N., 2022. Rustworkx: A high-performance graph library for Python. *J. Open Source Softw.* 7, 3968:1–3968:32.
- Turner, D.A., 1985. Miranda: A non-strict functional language with polymorphic types. In: Jouannaud, J.P. (Ed.), *Proceedings of the 1st International Conference on Functional Programming Languages and Computer Architecture. FPCA'85*, Springer, Nancy, France, pp. 1–16.
- Velez, M., Jamshidi, P., Sattler, F., Siegmund, N., Apel, S., Kästner, C., 2020. ConfigCrusher: Towards white-box performance analysis for configurable systems. *Autom. Softw. Eng.* 27, 265–300.
- Velez, M., Jamshidi, P., Siegmund, N., Apel, S., Kästner, C., 2021. White-box analysis over machine learning: Modeling performance of configurable systems. In: *Proceedings of the 43rd International Conference on Software Engineering. ICSE'21*, IEEE, Madrid, Spain, pp. 1072–1084.
- Velez, M., Jamshidi, P., Siegmund, N., Apel, S., Kästner, C., 2022. On debugging the performance of configurable software systems: Developer needs and tailored tool support. In: Damian, D., Zeller, A. (Eds.), *Proceedings of the 44th International Conference on Software Engineering. ICSE'22*, ACM, Pittsburgh, USA, pp. 1571–1583.
- Venkatakeerthy, S., Aggarwal, R., Jain, S., Desarkar, M.S., Updrasta, R., 2020. IR2Vec: LLVM IR based scalable program embeddings. *ACM Trans. Archit. Code Optim.* 17, 32:1–32:27.
- Wohlin, C., Runeson, P., Höst, M., Ohlsson, M.C., Regnell, B., Wesslén, A., 2012. *Experimentation in Software Engineering*. Springer.



Federico Bruzzone is currently a Ph.D. student in Computer Science at Università degli Studi di Milano, Italy. He was born in 2000 and since he was a child he has been passionate about computer science and music. He got his bachelor degree in Musical Computer Science, the master degree in Computer Science and currently he is involved in the research activity of the ADAPT Lab. His main research interests are (but are not limited to) programming languages and compilers, software maintenance and evolution. For any question he can be contacted at federico.bruzzone@unimi.it.



Walter Cazzola is currently a Full Professor in the Department of Computer Science of the Università degli Studi di Milano, Italy and the Chair of the ADAPT laboratory. Dr. Cazzola designed the mChARM framework, @Java, [a]C#, Blueprint programming languages and he is currently involved in the designing and development of the Neverlang language workbench. He also designed the JavAdaptor dynamic software updating framework and its front-end FiGA. He has written over 100 scientific papers. His research interests include (but are not limited to) software maintenance, evolution and comprehension, programming methodologies and languages. He served on the program committees or editorial boards of the most important conferences and journals about his research topics. He is associate editor for the *Journal of Computer Languages* published by Elsevier. More information about Dr. Cazzola and all his publications are available at <https://cazzola.di.unimi.it> and he can be contacted at cazzola@di.unimi.it for any question.



Luca Favini, is a Master's student in Computer Science at the Università degli Studi di Milano, Italy, where he also earned his Bachelor's degree in 2025. With a background in competitive programming, his interests range from theoretical aspects, specifically functional programming languages and algorithmic game theory, to practical applications in systems infrastructure. He is currently exploring these diverse areas to define his research specialization. He can be contacted at luca.favini1@studenti.unimi.it.