


A language for explaining counterexamples

Ezequiel José Veloso Ferreira Moreira ✉ 

Universidade do Minho, Braga, Portugal

INESC TEC, Braga, Portugal

José Creissac Campos ✉ 

Universidade do Minho, Braga, Portugal

INESC TEC, Braga, Portugal

Abstract

Model checkers can automatically verify a system's behavior against temporal logic properties. However, analyzing the counterexamples produced in case of failure is still a manual process that requires both technical and domain knowledge. However, this step is crucial to understand the flaws of the system being verified. This paper presents a language created to support the generation of natural language explanations of counterexamples produced by a model checker. The language supports querying the properties and counterexamples to generate the explanations. The paper explains the language components and how they can be used to produce explanations.

2012 ACM Subject Classification Theory of computation → Formal languages and automata theory

Keywords and phrases Model Checking, NuSMV, counterexample, natural language explanation

Digital Object Identifier 10.4230/OASICS.SLATE.2024.11

Funding *Ezequiel José Veloso Ferreira Moreira*: acknowledges FCT (Fundação para a Ciência e Tecnologia) grant 2023.01639.BD, funded by the ESF (European Social Fund) and PQDI (*Programa Demografia, Qualificações e Inclusão*).

1 Introduction

It is often important to guarantee that a given system behaves correctly under all possible operating conditions. One way to verify if operating requirements are met is through Model Checking [6]. This requires a formal representation of the system (a specification) and expressing requirements as temporal logic formulas (the properties). A model checker can then automatically verify the specification's behavior against the properties. Should the verification fail, the model checker will attempt to produce counterexamples. That is, behaviours of the specification that do not satisfy the properties being verified. Understanding these counterexamples is fundamental to understanding the failures of the specification and, by extension, of the system being specified.

However, the interpretation of counterexamples can be complex and time consuming, representing a barrier to Model Checking adoption. With the aim of simplifying this interpretation process, we are developing an approach to generate natural language explanations for counterexamples. The approach makes use of property specification patterns [8, 10] to identify templates of natural language explanations for the counterexamples produced. The templates need to query both the property and the counterexample to determine relevant information to include in the explanation. To that effect, a pattern-matching language was defined. This paper describes this language and how it can be used to produce explanations.


The paper is structured as follows: Section 2 presents the state of the art, Section 3 discusses Model Checking, focusing on temporal logic and the counterexamples produced, Section 4 introduces the explanation methodology, Section 5 explains the language that was created to support it, Section 6 present an example of usage, and Section 7 presents conclusions and possible future developments.



© Ezequiel José Veloso Ferreira Moreira and José C. Campos;
licensed under Creative Commons License CC-BY 4.0

13th Symposium on Languages, Applications and Technologies (SLATE 2024).

Editors: Mário Rodrigues, José Paulo Leal, and Filipe Portela; Article No. 11; pp. 11:1–11:14



Open Access Series in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 **2 State of the art**

46 Several approaches have been proposed to support the process of understanding Model
47 Checking results. Kaleeswaran et al. [12] provide an overview of the state of the art on the
48 topic, identifying the need to support non-experts in understanding model checking results
49 and natural language explanations as a promising direction for research.

50 Several approaches aim to produce explanations in natural language, which can more
51 easily be understood even by non-experts in Model Checking. Van den Berg et al. [13], for
52 example, propose a tool to interpret a counter-example and produce a description of the
53 error that caused the safety violation. Their work is aimed at the railway domain.

54 ASSERT [7] is another tool that produces explanations. It uses ACL2s to analyze a
55 series of requirements stored in an ontology. Should there exist a conflict between the
56 requirements, the counterexample produced by ACL2s is used to produce a controlled natural
57 language explanation that points to the requirements in conflict, the variables involved in
58 said requirements, the type of error detected, and the state that led to the conflict.

59 Another example is AnaCon [1], which analyzes normative texts to determine conflicts
60 within said texts. This is done first by writing these texts in a controlled natural language
61 form, which is then converted into the formal language CL. This formal representation of the
62 normative texts is then analyzed using the CLAN tool, which will produce a counterexample
63 should a conflict exist. This counterexample is then converted into the same controlled
64 natural language in which the normative text was first written, thus producing an explanation
65 for the counterexample.

66 Yet another example of this is presented by Lu Feng et al. [9], which produces a series of
67 structured language sentences to explain violations detected in a robot mission plan. This is
68 done by first specifying the path taken by the robot as a Markov decision process (MDP) and
69 analyzing this MDP formally using the PRISM model checker. Should a violation occur, then
70 a counterexample is produced, which is itself an MDP. The counterexample is then explained
71 by using controlled natural language sentences, recreating step-by-step the movement of the
72 robot that leads to the violation.

73 **3 Formal Methods and Model Checking**

74 Formal Methods focus on the specification, verification, and implementation of computing
75 systems through rigorous mathematical methods [14]. These methods are typically used
76 within safety and mission-critical areas, such as railways, avionics, and finance [11].

77 Model Checking is a Formal Methods technique that focuses on the formal verification
78 of a system's specification to determine if its behavior satisfies a series of properties [5].
79 These properties are defined in a given temporal logic (in this paper, CTL [4]), and then the
80 specification is verified using a model checker (in this paper, NuSMV [3]). If this verification
81 fails (i.e., the specification's behavior does not satisfy the property), the model checker will
82 attempt to produce a counterexample. That is, a sequence of states (a path) in the system's
83 specification that violates the property being verified.

84 **3.1 CTL properties**

85 Model checkers typically work with finite state representations of systems (e.g., Kripke
86 structures). These consist of a finite set of states and transitions between these states,
87 capturing the possible behaviors of the system. States are decorated with attributes that

88 help describe the system state at each moment. A path corresponds to a sequence of states
89 (with their attributes' values) representing a possible behavior of the model.

90 CTL (Computational Tree Logic) is a branching-time temporal logic that is used to
91 express and reason about the behavior of a system. Unlike linear time logic, which considers
92 a single timeline, CTL allows multiple potential futures (paths) to be considered at any given
93 state. For this purpose, the logic has both quantifiers over paths and temporal operators
94 combined into pairs. First, a path quantifier specifies the scope of the paths in which the
95 property must be verified. Then a temporal operator specifies in which states of any given
96 path the property must be verified.

97 The available path quantifiers are the operators **A** (all paths) and **E** (exists a path). If
98 operator **A** is used, the subsequent property must be verified in every future path in the
99 model that starts in the current state. If operator **E** is used, then at least one path must
100 exist, starting from the current state, that verifies the subsequent property.

101 Regarding temporal operators, CTL supports the following operators: **G** (globally), **F**
102 (finally), **X** (next), and **U** (until). If the **G** operator is used, then the property must hold in
103 every state of the path; if the **F** operator is used, then the property must eventually hold
104 at some point in the path; if the **X** operator is used, then the property must hold in the
105 next state in the path. The **U** operator expresses that one property must hold until another
106 property becomes true.

107 Pairing the path and temporal operators allows the expression of complex properties
108 over the system's behavior. For example, the property **AG** ϕ means that ϕ must hold in
109 every state of every possible path in the future, while if **EF** ψ is used, then there must exist
110 at least one path that in a future state verifies ψ . Similarly, **A**[ϕ **U** ψ] checks that ϕ hold
111 until ψ hold in every path, while **E**[ϕ **U** ψ] requires the existence of at least one path that
112 verifies ϕ until ψ . The properties ϕ and ψ above can themselves be CTL properties, or simple
113 propositional logic expressions over the attributes of the state.


114 Simple propositional logic expressions are atomic formulas expressed in propositional
115 logic with equality. Besides the usual propositional connectives (not \neg , and $\&$, or \vee ,
116 implies \rightarrow , iff \leftrightarrow , *xor* and *xnor*), and equality and inequality operators ($=$, \neq , $>$, $<$,
117 \geq , \leq), the language supports also testing for set membership (*in* connective). Using these
118 connectives, we can now write properties such as **AG**((*power* $>$ 10) $\&$ (*weight* $<$ 9)) that
119 expresses that in every state of every path the attribute *power* should have a value above 10
120 and the attribute *weight* a value below 9, or the property **EG**((*power* = 2) \rightarrow (*weight* = 8))
121 that expresses that there exists a path where it is always true that the attribute *power*
122 having a value of 2 implies that the attribute *weight* has a value of 8.

123 3.1.1 Counterexamples

124 As stated above, when the specification exhibits a behavior that does not satisfy the CTL
125 property, the model checker attempts to produce a counterexample. Useful counterexamples
126 are typically generated for safety properties. That is, properties expressing that undesirable
127 states will not be reached or undesirable behaviors are not possible in the system. In those
128 cases, the counterexample consists of a path that shows a behavior of the system that violates
129 the property being verified (the undesirable state being reached or the behavior happening).
130 It is important to note that while the counterexample contains a finite number of states, it
131 does not necessarily have a finite length in that a counterexample can end in a looping suffix.

132 In the case of liveness properties, that is, properties that express that desirable states of
133 behaviors are possible, generating useful counterexamples is harder, as the problem is the
134 system's lack of appropriate behavior.

11:4 A language for explaining counterexamples



■ **Figure 1** The state diagram that represents the ice cream machine’s formal model.

135 Interpreting the produced counterexamples is a crucial aspect of performing model
 136 checking. While the verification itself is automated, the counterexample analysis must be
 137 done manually. The goal is to determine the cause of the problem. This entails, first,
 138 identifying, in the counterexample, what the problematic states or behaviors are, and second,
 139 why (or whether) these states or behaviors are problematic. The conclusion might be that the
 140 system’s specification needs to be corrected, or it might be that the temporal logic property
 141 needs to be refined to better capture the intended requirement. Clearly, this analysis requires
 142 an understanding of both the system’s specification and the system’s domain.

143 3.2 An example

144 A simple example of a system is now presented to allow for greater clarity on the various
 145 concepts being discussed. The example used is a washing machine that can wash, rinse, and
 146 spin laundry, with various programs that determine how long each washing phase takes. To
 147 support this, the machine’s specification defines a number of state attributes, including the
 148 current phase and its duration, the current state of the machine (starting a phase, ..., etc.),
 149 and an error code in case of malfunctioning.

150 We also define as one of its design requirements that *The machine must always either*
 151 *have its internal error code set to 0 or the machine’s current state must be the error state.*
 152 Figure 1 presents the machine’s control logic, focusing on the phase attribute.

153 To use a model checker (in this case, NuSMV) to verify this requirement, the first step is
 154 to specify the system formally. This was achieved using the IVY workbench tool [2], but
 155 discussing the model is outside this paper’s scope. Herein, we are specifically interested in
 156 the counterexamples produced. It is enough to say that it represents the state machine in
 157 Figure 1.

158 With the model created, the requirement was converted into the CTL property:

$$159 \quad \mathbf{AG}(error = 0 \mid state = errorState). \quad (1)$$

160 Attempting to prove this property fails, and the counterexample presented in Listing 1 is
 161 produced. The challenge now is to generate an explanation of the trace in natural language.

162 4 Explanation Methodology

163 The methodology aims to produce explanations for counterexamples generated by a model
 164 checker to simplify their interpretation. To achieve this, it uses three different sources of
 165 information (inputs) to generate the explanations: the counterexamples produced by the
 166 model checker (and corresponding properties), an explanation templates library, and domain
 167 information.

■ **Listing 1** An example of a CTL property and resulting counterexample (in the required format to generate an explanation). Note that if an attribute has the same value between sequential states then that value is not shown, for size and readability reasons.

```

AG (error = 0 | state = errorState)      ---
--> State: 31.1 <-
    duration = none
    phase = standby
    door = open
    load = empty
    bleach = empty
    state = startingPhase
    program = 0
    error = 0
    water = empty
    action = nil
-> State: 31.2 <-
    load = overloaded
    state = errorState
    error = 4
    action = excessiveLoadDetected
-> State: 31.3 <-
    load = loaded
    state = beginDoorCheck
    action = sufficientLoadRemoved

```

168 The process of producing counterexamples starts with identifying a suitable explanation
 169 template consisting of partially filled-in sentences that will be completed with information
 170 from the various inputs to generate an explanation. In this section, we explain each of the
 171 inputs to the process and how they are used. A tool that implements the proposal is being
 172 developed and mentioned where deemed relevant.

173 4.1 Verification results input

174 The verification results input is used to provide the data related to the counterexample that
 175 needs to be explained.

176 Each item within this input consists of a pair (**property,counterexample**), where
 177 the **property** is the CTL property whose verification was attempted and failed, and the
 178 **counterexample** is the counterexample produced by the model checker for that **property**.
 179 An example of such an item can be found in Listing 1.

180 4.2 Explanation templates library

181 The explanation templates library provides the templates to produce the explanations for
 182 each property/counterexample pair. Listing 2 presents three examples of entries in this
 183 library.

184 The library contains a series of entries that consist of tuples (**pattern, condition-**
 185 **alTemplate1, ..., conditionalTemplateN**), where the **pattern** is used to dictate the
 186 types of **property** that can be explained using the entry, and each **conditionalTemplate**
 187 corresponds to a pair (**condition, template**). The **condition** is a logical expression over
 188 the counterexample that must hold for the corresponding **template** to be used for the

11:6 A language for explaining counterexamples

■ **Listing 2** An example of three dictionary entries, with patterns that the property in Listing 1 matches

```
AG({E}{1,eq,x,y} | {E}{2,eq,x,y}) ----
Is should always be the case that either {DSE}{-,{E}{1}},
or {DSE}{-,{E}{2}}.
However, in state {S}{1,invert,sv} the two conditions do not hold,
with the verified values of {E}{1,-,x} and {E}{2,-,x} being
{S}{1,invert,t,{E}{1,-,x}} and
{S}{1,invert,t,{E}{2,-,x}}, respectively. ;;;;
AG({E}{1} | {E}{2}) ----
It was expected that in every state of the machine
either {DSE}{-,{E}{1}} or {DSE}{-,{E}{2}}.\n
However on state {S}{1,invert,sv} neither
{DSE}{-,{E}{1}} nor {DSE}{-,{E}{2}}. ;;;;
AG({0}{1,or}) ----
It was expected that in every state of the machine
{DSE}{-,{0}{1}} was always verified.\n
However at the end of the path {C}{2,3},
{DSE}{negateLeftSide,{0}{1}} ;;;;
```

189 explanation. It consists of a conjunction of sub-conditions expressed over the states of the
190 counterexample.

191 A **template** will be used to explain the counterexample for a property if the **property**
192 matches the **pattern** and the **condition** holds over the counterexample. As seen in Listing 2,
193 a template consists of a mix of natural language and expressions that will query the property
194 and counterexample for information. The language used to write these expressions (also used
195 in the conditions above) will be discussed in the next section.

196 A special case occurs when there is only one **template** to be used if the **pattern** matches
197 the property. For this specific case, the item becomes a simple pair (**pattern,template**)
198 within the library. The three examples in Listing 2 fall in this category.

199 Should the data item not be matched by any **pattern**, or fail to meet all of the **conditions**,
200 then the explanation can not be created, and the message given to the user will be *No*
201 *matching pattern found for the specified property.*

202 4.3 Domain information input

203 This input is used to provide additional domain context that can be utilized during the ex-
204 planation process. Examples of domain information entries are shown in Listing 3. They take
205 the form of tuples (**domain pattern, contextExplanation1, ..., contextExplanationN**).
206 The **domain pattern** identifies which expression in the specification can be explained using
207 the particular entry. Each **contextExplanation** consists of a pair (**context,explanation**),
208 where **context** is used to identify the context that the **domain pattern** is used in, and
209 **explanation** is the explanation to be used in that context.

210 The notion of context enables control over how an expression captured by the domain
211 pattern may be explained. Different contexts represent different ways to explain the same
212 expression, depending on where they appear in the template. This enriches the explanations
213 that can be produced, as it allows expressions to be described differently in different parts
214 of the explanation. For example, differentiating between positive and negative contexts, or
215 using long and short versions of the description to avoid repetition.

■ **Listing 3** An example of four domain items

```

error = 0 ---- the current error code is 0, meaning no error
           is currently occurring ;;;;
error = {V}{x} ---- the current error code is {V}{x} ;;;;
state = errorState ---- the current state is the error state;
error = 0 | state = errorState
\\ \\ default ---- either the state is the error
           state, or the current error code is 0.
\\ \\ negateLeftSide ---- the current error code is not 0,
           and despite this the current state is the error state
\\ \\ negateRightSide ---- the current state is the errorState,
           and despite this current error code is 0
; ; ; ;

```

216 A special case for these items occurs when there is only one **context** with the value
 217 **default**. In this case, the item becomes a pair (**domain pattern, explanation**) within the
 218 domain input.

219 5 The language

220 As is clear from the discussion above, to instantiate the explanation templates, we need a
 221 means to query the CTL properties and the resulting counterexamples. To achieve this,
 222 we have defined a language that allows us, through pattern matching, to identify relevant
 223 elements in the verification results.

224 Several requirements for this language can be identified from the discussion. Regarding
 225 the writing of CTL patterns in the explanation templates library:

- 226 1. The language to be defined must support writing **patterns** to match properties written
 227 in the CTL language, as described in Section 3.1.
- 228 2. The tokens in the language, however, must only match propositional logic expressions.
 229 Temporal and path operators must be included explicitly in the patterns, not matched.
 230 This is because knowing the structure of the temporal operators in the property is relevant
 231 to deciding which type of explanation to provide. Allowing tokens in the language to
 232 match temporal sub-expressions would imply adding recursion to the natural language
 233 templates. While this would add expressive power, we feel the added complexity (both in
 234 terms of the writing of the explanations templates library and of its processing and use)
 235 is not justified.

236 Regarding the natural language templates:

- 237 3. The tokens in the language must allow for the explanation **template** to reference
 238 information in the property/counterexample pair whose property was matched by its
 239 associated **pattern**.
 - 240 a. The tokens must allow reference to the **property**, specifically to any part of it that
 241 has been matched using another token.
 - 242 b. The tokens must allow reference to any part of the **counterexample**, specifically any
 243 sequence of states, singular state, or specific attributes of any given state.
- 244 4. The language must allow for the **condition** in the explanation templates library to be
 245 expressed using the information in the **counterexample**. The goal here is to support
 246 tailoring the explanation to the counterexample produced when a property can fail in
 247 several ways.

- 248 5. The language must support the usage of domain information to replace information
 249 captured from the property or counterexample in an explanation **template**. This will
 250 allow the natural language explanation to be less technical and expressed in terms closer
 251 to the domain.
- 252 6. Similarly to Requirement 3, the language must allow the explanation **template** to
 253 reference the information matched by tokens in a domain item's **pattern**.

254 To fulfill these requirements, the language includes a series of tokens. All of the tokens
 255 consist of two parts, each encapsulated within **{}**: the type of the token, which dictates its
 256 main purpose, and the parameters, which modify the token's behavior. The types defined
 257 are **E** (Expression), **O** (Operation), **S** (State), **C** (Counterexample path), **DSE** (Domain
 258 Specific Explanation) and **V** (Value). The first two can be used in both the patterns and
 259 templates of the explanation patterns library. The following three are used in the templates
 260 only. The final one is used in domain items (both domain patterns and explanations).

261 Different token types have different required and optional parameters, but the parameters
 262 attached to each token type are the same independently of where the token is used. However,
 263 how the parameters modify the behavior of the token can change depending on where the
 264 token is being used. Parameters' values can be provided by naming the parameters explicitly
 265 (e.g., $\{O\}\{identifier : 1\}$), if names are not provided, the order of the values becomes relevant.
 266 For the sake of brevity, we will employ the unnamed variant (in this case $\{O\}\{1\}$).

267 5.1 Tokens in patterns

268 Tokens of type **E** and **O** can be used in the explanation templates library patterns to help
 269 define the type of property that will use a particular explanation template.

270 5.1.1 E tokens in patterns

271 The **E** token matches simple expressions only, with the option of further restricting what
 272 type of expression can be matched. This type of token has one required parameter, the
 273 **identifier**, as well as three optional parameters: **operation**, **nameV1**, and **nameV2**.

274 The **identifier** is a positive integer used to uniquely identify a token of this type, such
 275 that two **E** tokens with the same identifier must refer to the same simple expression. Thus, if
 276 the same **identifier** is used in two **E** tokens in a **pattern**, then the **property** being matched
 277 must have the same expression in both places, or else it will not match the **pattern**.

278 The **operation** parameter defines the type of simple expression the token will match.
 279 The values for this parameter are related to the available operators as follows: **eq** relates to
 280 **=**, **diff** relates to **!=**, **gt** relates to **>**, **lt** relates to **<**, **gte** relates to **>=**, **lte** relates to **<=**,
 281 **in** relates to **in** and **-** relates to any operation, serving as a wildcard for this parameter.

282 **nameV1** and **nameV2** are both identifiers used to internally reference the values to the
 283 left and right of the operator, respectively, storing the values so that they may be explicitly
 284 used later in the explanation template. This allows for the explicit reference to each side of
 285 the expression and is necessary to fulfill Requirement 5.

286 Some examples of these tokens' use can be seen in Listing 2, such as with the second
 287 example's $\{E\}\{1\}$ or the first example's $\{E\}\{2, eq, x, y\}$. In the first case, the token is used
 288 to match any simple expression and will then be identified using the identifier 1. When the
 289 same token is used in the **template**, its value will be whatever was matched by the token in
 290 the CTL property. In the second case, the token will only match an equality expression, will
 291 use the identifier 2, and will store the left and right sides of the expression in the variables x
 292 and y respectively, such that their values can be accessed in the **template** explicitly.

293 With these options, the token can be used in the dictionary item's **pattern** to specify
 294 any type of simple expression that can be matched, partially fulfilling Requirement 1.

295 5.1.2 O tokens in patterns

296 The **O** token matches expression with propositional connectives only, again with the option
 297 of further restricting which connective can be matched. This token type has one required
 298 parameter, the **identifier**, and one optional parameter, the **operation**.

299 The **identifier** works as for the **E** token, being a positive integer that uniquely identifies
 300 the token. In the same manner, a **pattern** with two **O** tokens with the same identifier will
 301 only match **properties** that have the same exact expression in both places.

302 The meaning of **operation** differs, as it refers to the propositional connectives that the
 303 token can match. The values for this parameter are related to the connectives as follows: **not**
 304 relates to **!**, **and** relates to **&**, **or** relates to **|**, **imp** relates to **->**, **equi** relates to **<->**, **xor**
 305 relates to **xor**, **xnor** relates to **xnor** and **-** relates to any operation, serving as a wildcard for
 306 the parameter.

307 An example of this token's use can be seen in Listing 2, in the CTL property of the third
 308 example: $\{O\}\{1, or\}$. This token is used to match expressions with the connective **|**.

309 As illustrated, the token supports matching any propositional connective in the patterns
 310 of the explanation templates library's entries. This, in combination with the previous **E**
 311 token, allows for the fulfillment of Requirement 1. Since neither **O** type tokens nor **E** type
 312 tokens match temporal or path operators, Requirement 2 is also fulfilled.

313 5.2 Tokens in templates

314 The tokens that can be used in an explanation templates library entry's **template** are **E**, **O**,
 315 **S**, **C**, and **DSE**. Each of them represents a different part of the information captured in the
 316 verification results' pair to be added to the explanation.

317 5.2.1 E tokens in templates

318 Tokens of type **E** in a **template** serve to identify a simple expression that was matched in
 319 the **pattern** in order to use the captured information. The **identifier** parameter is used to
 320 determine which of the matched expressions the token refers to, such that an **E** token in the
 321 **template** will always refer to an **E** token in the **pattern** that has the same **identifier**.

322 Both the **nameV1** and **nameV2** parameters may be used in the **template** to refer to a
 323 specific side of the expression, provided that those same options exist in the corresponding **E**
 324 token that was used in the **pattern**. However, only a single one of these can be used per **E**
 325 token in the **template**, and if neither is specified, then the **E** token will refer to the complete
 326 simple expression. This occurs because the only context where it makes sense to use both
 327 sides is when talking about the whole expression.

328 The **operation** parameter is not relevant when the token is used in a **template**. The
 329 matching is done on the CTL property. Here, the goal is to refer to and use what has been
 330 matched. Hence, if used, the value of the **operaton** parameter will be ignored.

331 Some examples of these tokens can be seen in Listing 2, such as in the first example's
 332 template: $\{E\}\{1, -, x\}$ and $\{E\}\{1, -, y\}$. For both of these cases, the expression that is
 333 being referenced is the token $\{E\}\{1, eq, x, y\}$ in the **pattern**. In the first case, the expression
 334 will be replaced with the value stored in the x variable. In contrast, in the second case, it
 335 will be replaced with the value stored in the y variable, corresponding to the left and right
 336 sides of the matched expression, respectively.

11:10 A language for explaining counterexamples

337 With these options, the token can partially use any matched simple expression in the
338 **template**, allowing it to partially fulfill Requirement 3a.

339 5.2.2 O tokens in templates

340 Tokens of type **O** in a **template** serve to identify a logical operation that was matched
341 in the **pattern**. As before, the **identifier** parameter is used to determine which of the
342 matched logical operations the token refers to, the same way as with the **E** token. Similarly,
343 the **operation** parameter does nothing when used in the **template**, and its value will be
344 ignored.

345 An example of these token's use can be seen in Listing 2, in the third example's template:
346 $\{O\}\{1\}$. This token is used to refer to the **pattern** token $\{O\}\{1, or\}$, and will be replaced
347 in the template with the value of the matched operation.

348 With these parameters, the token can use any logical operation in the **template**, as long
349 as they can be matched in the CTL property. Combined with the previous **E** token, this
350 allows for the complete fulfillment of Requirement 3a.

351 5.2.3 S tokens

352 Tokens of type **S** serve to query the **counterexample** for information regarding a specific
353 state. This type of token has three required parameters, **index**, **order**, and **type**, as well as
354 two optional parameters, **token** and **condition**.

355 The **type** parameter is used to distinguish what type of information the token will be
356 replaced with in the template. If its value is **t**, then the replacement will use the state's
357 attributes, while if the value is **sv**, then only the state's identifier will be utilized.

358 The **token** parameter can have any attribute name as its value and, when defined, will
359 restrict the replacement made to only the value of the given attribute in the state. For this
360 reason, this parameter only takes effect when the type parameter has value **t**.

361 The remaining parameters are used to identify which state the token refers to. In its
362 simplest form, we can use the **index** and **order** parameters. **index** is a positive integer that
363 acts as an index over the trace. **order** can be either **default** or **invert** and defines which
364 end of the trace we want to count from. If the **order** parameter has value **default**, then the
365 state the token is referencing will be the **index**th one of the counterexample. If the **order**
366 parameter has value **invert**, then the state the token is referencing will be the **index**th last
367 one of the counterexample.

368 Always counting states from the start or the end of the counterexample is too restrictive,
369 so it is also possible to do it from a state that satisfies a given condition. To do this, we
370 must use the **condition** parameter. This condition is defined as a conjunction of simple
371 expressions and **E** tokens. When it is defined, then the first state of the counterexample where
372 the expression becomes true is used as a reference, and the state the token is referencing be
373 the **index**th one before or after it, depending on whether the **order** parameter is **default** or
374 **invert**, respectively.

375 Several examples of these tokens can be seen in Listing 2. Two examples are the
376 $\{S\}\{1, invert, sv\}$ and $\{S\}\{1, invert, t, \{E\}\{2, -, x\}\}$ expressions in the first item's template.
377 In the first case, the token will be replaced using the state identifier for the first state that
378 occurs counting from the end, i.e., with the identifier of the last state of the counterexample.
379 In the second case, the token will be replaced with the value of the attribute represented by
380 $\{E\}\{2, -, x\}$ (the left side of the expression $\{E\}\{2, eq, x, y\}$ that is captured in the **pattern**)
381 on the last state on the counterexample.

382 With these options, the token can be used to reference states in the **counterexample**
 383 and obtain all its relevant data, fulfilling the part of Requirement 3b related to single states
 384 and specific attributes in any given state.

385 5.2.4 C tokens

386 Tokens of type **C** match sequences of states, i.e., paths, in the **counterexample**. This type
 387 of token has one required parameter, **start**, and two optional parameters, **end** and **token**.

388 The start **start** parameter defines the state where the path starts. It can have multiple
 389 different types of values, which affect the behavior of the token. If this parameter is a
 390 positive integer, then the start of the path will be that state in the **counterexample**. If this
 391 parameter is an **E** token, then the path will start at the first state of the **counterexample**
 392 where the simple expression the token represents first becomes true. If this parameter is the
 393 word **loop**, then the path will start at the first state of the loop within the **counterexample**.
 394 If this parameter is the word **full**, then the path will be the entire **counterexample**.

395 The **end** parameter defines the end of the path the token represents. It is always a
 396 positive integer, and if the **start** parameter is not **full**, then the path will end at the state
 397 indicated by this option. If the **start** parameter has value **full**, then the **end** parameter has
 398 no meaning.

399 The **token** parameter takes as value an attribute. If set, the parameter will change the
 400 replacement process for this token in the **template** so that only the value of the identified
 401 attribute will be present in each state of the path.

402 An example of this token's use can be seen in Listing 2, in the third example's template:
 403 $\{C\}\{2,3\}$. In this example, the token will be replaced by the path starting on the second
 404 state of the counterexample and ending on the third, with every state showing the values of
 405 all state attributes.

406 With these options, the token can be used to reference any sequence of states in the
 407 **counterexample** and obtain the relevant data, and in combination with the **S** token fulfilling
 408 Requirement 3b.

409 5.2.5 DSE tokens

410 Tokens of type **DSE** serve to replace parts of information in the **template** using their
 411 counterpart domain information. This type of token has two required parameters, **value** and
 412 **context**.

413 The **value** parameter can take as value an attribute, the value of an attribute, a simple
 414 expression, a propositional connective, an **E** token, or an **O** token. Its value will then be
 415 matched against a domain item's **domain pattern**, which will then be used to replace the
 416 token within the **template**.

417 The **context** parameter specifies the context that should be used when replacing the
 418 token in the template with a value from a context item of the domain information input. Its
 419 value is either an attribute or -. In the latter case, the value of the context becomes **default**.

420 Some examples of these tokens can be seen in the templates of Listing 2, such as with
 421 the first item's $\{DSE\}\{-, \{E\}\{1\}\}$ and the third item's $\{DSE\}\{negateLeftSide, \{O\}\{1\}\}$.
 422 For the first case, the token will be replaced by the domain item's information regarding the
 423 simple expression that is in the token $\{E\}\{1\}$ (i.e., the left side of the **or** expression in the
 424 CTL property). For the second case, the same will occur, using the **or** expression captured
 425 by $\{O\}\{1\}$, but instead of using the information in the **default** context, the explanation
 426 attached to the **negateLeftSide** context will be used instead.

11:12 A language for explaining counterexamples

427 With these options, the token can use domain information to replace any token being
428 used in a template within the explanations templates library, thus fulfilling Requirement 5.

429 5.3 Tokens in the domain information

430 Only one type of token exists within the domain information, the **V** token. This token allows
431 capturing values in a **domain pattern** and then using those values in the corresponding
432 explanations. This type of token has one required parameter, **value**, and an optional
433 parameter, **joinOperation**.

434 The **value** parameter is an identifier that is used to distinguish the data that was matched
435 in the **domain pattern**, such that a **V** token with a different identifier must have matched
436 a different value and vice-versa. This also means that if two **V** tokens have the same **value**
437 in the **domain pattern**, then they must have matched the same value in multiple parts of
438 said **domain pattern**.

439 The **joinOperation** parameter can either take the value **and** or the value **or** and is
440 used to replace a list of values matched in the **domain pattern** with the enumeration of
441 the values expressed as conjunction or disjunction, respectively. To be more precise, if the
442 join operation is **and** and the list is $[m_1, m_2, \dots, m_{n-1}, m_n]$, then the value will be replaced by
443 " m_1, m_2, \dots, m_{n-1} and m_n ", with something similar occurring when using the **or** option.

444 An example of this token can be seen in Listing 3, in the second example: $\{V\}\{x\}$. In
445 this example, the pattern will capture any simple expression whose left side has the value
446 **error**, and the token will store the corresponding right side with the identifier x , allowing
447 for the usage of this value in the corresponding **explanation** by using the same token.

448 The token is able to use information captured within the **domain pattern** in the paired
449 **explanation**. This explanation will then be used to replace a **DSE** token in a **template**,
450 thus fulfilling Requirement 6.

451 6 Using the language

452 With the language presented, we can finally showcase its use with the example presented in
453 Section 3.2. A tool is being developed to support the process briefly described in Section 1.
454 The tool is already able to read the verification results, an explanation templates library,
455 and a domain information file and produce explanations for the counterexamples using the
456 templates in the library.

457 To apply it to the example in Section 3.2, the property, as well as the counterexample
458 that resulted from its verification attempt, were used as input to the tool (see Listing 1). To
459 generate the explanations, an explanations templates library (see Listings 2 for an excerpt)
460 and the domain input file in Listing 3 were used. With all the inputs available, the tool
461 selected the first item of the dictionary that matched the **property**. In this case, the first
462 item shown was selected.

463 The explanation was generated using this item by filling in the tokens within the **template**,
464 according to the previously presented language. As such, the token $\{DSE\}\{-, \{E\}\{1\}\}$
465 was replaced by the domain information in the **default** context relating to the expression
466 captured by the **pattern**'s token $\{E\}\{1, eq, x, y\}$, which corresponds to $error = 0$. In the
467 domain input, the expression $error = 0$ only has a single context. Its value is *the current*
468 *error code is 0, meaning no error is currently occurring*, and thus this phrase replaced the
469 token $\{DSE\}\{-, \{E\}\{1\}\}$ in the final explanation. A similar process occurred for each of
470 the **DSE** tokens, each being replaced with domain information relevant to each of the **E**
471 tokens in the **pattern**.

472 Meanwhile, the **E** tokens are filled in with the information of the captured expression
 473 corresponding to the token's **identifier** parameter. In the case of $\{E\}\{1, -, x\}$ its value will be
 474 replaced with the value associated with the x variable in the **pattern**'s token $\{E\}\{1, eq, x, y\}$,
 475 which is the left side of the capture expression that has the value of *error*. A similar process
 476 occurs for the other **E** token, but referring to the **pattern**'s token $\{E\}\{2, eq, x, y\}$ instead.

477 The last type of tokens present in this example are the **S** tokens. For the first token, using
 478 the *sv* value for the type parameter leads to the token being replaced using the identifier of
 479 the corresponding state. This state is identified by the other two parameters and is the first
 480 to last state of the counterexample. Thus, the value that replaces this token is the identifier
 481 of the last state of the counterexample: 31.3.

482 For the other **S** tokens, the replacement uses the **type** t , meaning the replacement will
 483 be done with the values of attributes in the specified state (the last state). Additionally,
 484 the **token** parameter has an **E** token specified, corresponding to the left side of one of the
 485 captured expressions, an attribute. Thus, the value used in the replacement is the value of
 486 the attribute specified by the **E** token: *error* in the first case (with value 4) and *state* in the
 487 second (with value *beginDoorCheck*).

488 The explanation created for the example given using the shown inputs was finally:

489 It should always be the case that either the current error code is 0, meaning no error
 490 is currently occurring, or the current state is the error state.

491 However, in state 31.3 the two conditions do not hold, with the verified values of error
 492 and state being 4 and *beginDoorCheck*, respectively.

493 **7 Conclusion and future work**

494 A model checker can automatically verify a system's behavior against temporal logic properties.
 495 However, in case of failure, analyzing the results is still a manual process that requires both
 496 technical and domain knowledge. To aid this process, we have been exploring the generation
 497 of natural language explanations for model-checking counterexamples.

498 The approach we are developing makes extensive use of CTL property patterns and natural
 499 language templates to generate explanations for counterexamples. This paper describes
 500 the language used to query properties and counterexamples to instantiate the explanation
 501 templates. In particular, we have explained the tokens that allow for complex templates to
 502 be created and the reasoning behind their definition.

503 The paper reports on ongoing work and several venues for future work can be identified
 504 and are being pursued. One obvious area of future work is the continued improvement
 505 of the language. This should be guided by requirements arising from the need to create
 506 quality explanation templates. We plan to develop a series of case studies to analyze the
 507 expressiveness of the approach and, as part of this, to carry out user studies to assess the
 508 quality of the produced explanations.

509 In any case, several improvements have already been identified that will improve the
 510 expressive power of the language. We plan to expand the functionality of the **S** token when
 511 there is a condition so that it can search beyond only the first state where the condition is
 512 verified or allow for verification of conditions that occur in a state before/after the search
 513 point. Additionally, we want to expand the **S** token's **condition** so that it can also search
 514 for logical operations referenced by a **O** token.

515 Other enhancements include adding the ability to reference states obtained in a previous
 516 token to the current one, such that it is possible to check a condition only once and then
 517 refer to the state that matches that condition in other future tokens of the template; the

518 ability for the **token** parameter in both the **S** and **C** tokens to contain more than a single
 519 attribute; the ability for the **C** token to be able to have an **S** token for both the **start** and
 520 **end** options; or the possibility to automatically change a **DSE** token's context in a template
 521 based on a given condition.

522 Finally, we plan to explore how generative AI might be integrated into the approach and
 523 whether that can improve the quality of the generated explanations.

524 ——— References ———

- 525 1 John J. Camilleri, Mohammad Reza Haghshenas, and Gerardo Schneider. A web-based tool for
 526 analysing normative documents in english. In *Proceedings of the 33rd Annual ACM Symposium*
 527 *on Applied Computing, SAC '18*, page 1865–1872, New York, NY, USA, 2018. Association for
 528 Computing Machinery. doi:10.1145/3167132.3167331.
- 529 2 J. C. Campos and M. D. Harrison. Interaction engineering using the ivy tool. In *ACM*
 530 *Symposium on Engineering Interactive Computing Systems (EICS 2009)*, pages 35–44, New
 531 York, NY, USA, 2009. ACM. doi:10.1145/1570433.1570442.
- 532 3 Alessandro Cimatti, Edmund Clarke, Enrico Giunchiglia, Fausto Giunchiglia, Marco Pistore,
 533 Marco Roveri, Roberto Sebastiani, and Armando Tacchella. Nusmv 2: An opensource tool for
 534 symbolic model checking. In *Computer Aided Verification*, pages 359–364, Berlin, Heidelberg,
 535 2002. Springer Berlin Heidelberg. doi:10.1007/3-540-45657-0_29.
- 536 4 E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent
 537 systems using temporal logic specifications. *ACM Transactions on Programming Languages*
 538 *and Systems*, 8(2):244–263, April 1986. doi:10.1145/5397.5399.
- 539 5 Edmund M Clarke. The birth of model checking. In Orna Grumberg and Helmut Veith,
 540 editors, *25 Years of Model Checking*, volume 5000 of *Lecture Notes in Computer Science*, pages
 541 1–26. Springer, 2008. doi:10.1007/978-3-540-69850-0_1.
- 542 6 Edmund M. Clarke, Jr., Orna Grumberg, and Doron A. Peled. *Model Checking*. MIT Press,
 543 Cambridge, MA, USA, 1999.
- 544 7 Andrew Crapo, Abha Moitra, Craig McMillan, and Daniel Russell. Requirements capture and
 545 analysis in assert(tm). In *2017 IEEE 25th International Requirements Engineering Conference*
 546 *(RE)*, pages 283–291, 2017. doi:10.1109/RE.2017.54.
- 547 8 M. B. Dwyer, G. S. Avrunin, and J. C. Corbett. Patterns in property specifications for
 548 finite-state verification. In *Proc. 21st International Conference on Software Engineering, ICSE*
 549 *'99*, pages 411–420. ACM, 1999. doi:10.1145/302405.302672.
- 550 9 Lu Feng, Mahsa Ghasemi, Kai-Wei Chang, and Ufuk Topcu. Counterexamples for robotic
 551 planning explained in structured language. In *2018 IEEE International Conference on Robotics*
 552 *and Automation (ICRA)*, pages 7292–7297. IEEE, 2018. doi:10.1109/ICRA.2018.8460945.
- 553 10 M.D. Harrison, P. Masci, and J.C. Campos. Verification templates for the analysis of user
 554 interface software design. *IEEE Transactions on Software Engineering*, 45(8):802–822, August
 555 2019. doi:10.1109/TSE.2018.2804939.
- 556 11 Anne Elisabeth Haxthausen. *An Introduction to Formal Methods for the Development of*
 557 *Safety-critical Applications*. DTU Orbit, 2010.
- 558 12 Arut Prakash Kaleeswaran, Arne Nordmann, Thomas Vogel, and Lars Grunske. A systematic
 559 literature review on counterexample explanation. *Information and Software Technology*,
 560 145:106800, 2022. doi:10.1016/j.infsof.2021.106800.
- 561 13 Lionel van den Berg, Paul Strooper, and Wendy Johnston. An automated approach for
 562 the interpretation of counter-examples. *Electronic Notes in Theoretical Computer Science*,
 563 174(4):19–35, 2007. Proceedings of the Workshop on Verification and Debugging (V&D 2006).
 564 doi:10.1016/j.entcs.2006.12.027.
- 565 14 Jim Woodcock, Peter Gorm Larsen, Juan Bicarregui, and John Fitzgerald. Formal methods:
 566 Practice and experience. *ACM Comput. Surv.*, 41(4), October 2009. URL: [https://doi.org/](https://doi.org/10.1145/1592434.1592436)
 567 10.1145/1592434.1592436.