# *Knowledge-Based Support for Adhesive Selection: Will it Stick?*

SIMON VANDEVELDE and JOOST VENNEKENS

*KU Leuven, De Nayer Campus, Dept. Of Computer Science,*
*J.-P. De Nayerlaan 5, 2560 Sint-Katelijne-Waver, Belgium*
*Leuven.AI — KU Leuven institute for AI, B-3000, Leuven, Belgium*
*Flanders Make — DTAI-FET*
(*e-mails:* s.vandevelde@kuleuven.be, joost.vennekens@kuleuven.be)

JEROEN JORDENS, BART VAN DONINCK and MAARTEN WITTERS

*Flanders Make, Oude Diestersebaan 133, 3920 Lommel, Belgium*
(*e-mails:* jeroen.jordens@flandersmake.be, bart.vandoninck@flandersmake.be,
maarten.witters@flandersmake.be)

## Abstract

As the popularity of adhesive joints in industry increases, so does the need for tools to support the process of selecting a suitable adhesive. While some such tools already exist, they are either too limited in scope or offer too little flexibility in use. This work presents a more advanced tool, that was developed together with a team of adhesive experts. We first extract the experts' knowledge about this domain and formalize it in a Knowledge Base (KB). The IDP-Z3 reasoning system can then be used to derive the necessary functionality from this KB. Together with a user-friendly interactive interface, this creates an easy-to-use tool capable of assisting the adhesive experts. To validate our approach, we performed user testing in the form of qualitative interviews. The experts are very positive about the tool, stating that, among others, it will help save time and find more suitable adhesives.

*KEYWORDS*: knowledge base, first order logic, decision support, adhesive selection

## 1 Introduction

The Flanders Make Joining & Materials Lab (FM JML) is specialized in adhesive bonding. They support companies in selecting the most appropriate adhesive for a specific use case, by accounting for characteristics such as strength, temperature resistances, adhesive durability, and more. Currently, this is done manually by one of the *adhesive experts* that work at the lab. Selecting a suitable adhesive is a time-consuming and labor intensive task, due to the large number of adhesives available on the market, each with extensive data sheets. Currently, the experts do not use any supporting tools to help them perform the selection, because the current generation of tools does not meet their requirements.

This paper describes our work on a logic-based tool which supports the experts in the selection process. It is structured as follows. We start by describing the process of

selecting an adhesive and the state-of-the-art tools in Section 2, and elaborate on the logical system used in this work in Section 3. Next, we present our Adhesive Selector Tool in Section 4, where we discuss the process of Knowledge Acquisition, how the system handles unknown parameter values, and how the experts interface with the knowledge. We share the results of our preliminary three-fold validation in Section 5 and the results of our comprehensive user study in Section 6. Finally, we describe our lessons learned in Section 7 and conclude in Section 8.

This paper is an extended version of (Vandevelde et al. 2022), as presented at the LPNMR22 conference. Its main addition is a user study with the target users of the JML lab, in the form of semi-structured interviews. Additionally, we elaborated on some sections and updated the text with various minor improvements.

## 2 Adhesive selection and current tools

As there is no universally applicable adhesive, the selection of an adhesive is an important process. There are many factors that influence the choice of an adhesive: structural requirements such as bonding strength and maximum elongation, environmental factors such as temperature and humidity, economic factors, and more. Due to the complexity of the problem, there is quite a potential for tools that support this selection process. Yet, Ewen (2010) concludes that "there is a severe shortage of selection software, which is perplexing especially when the task of adhesive selection is so important."

Currently, when tasked with a use case, the experts work in two steps. First, they try to identify requirements, such as temperature ranges or values for parameters like minimum strength. Based on this list of requirements, they perform an initial selection by manually looking through various data sheets while keeping track of which adhesives are suitable. In the second step, these initially chosen adhesives are put to the test by performing real-life experiments in FM's lab, to ensure suitability. However, this testing step is costly and time-consuming, so it is important that the initial selection is as precise as possible. While there are tools available for this process, the FM experts do not use them because they are either too simplistic, or not sufficiently flexible.

The most straightforward selection tools are websites offering simple interfaces[1] (da Silva et al. 2018). Based on a series of questions, they provide advice to support selection. However, they still require the expert to look up and process the information themselves.

There are also a number of expert systems to be found in the literature (Kannan and Prabu 2004; Allen and Vanderveldt 1995; Su et al. 1993; Moseley and Cartwright 1992; Meyler and Brescia 1993; Lees and Selby 1993; Lammel and Dilger 2002). Here, domain knowledge is captured and formalized in the form of rules, which can be used for adhesive selection by forward chaining and often also for generating explanations by backward chaining. However, these systems have a number of downsides: they are low in both interpretability and maintainability by the expert, often not all required

---

[1] Such as www.adhesivestoolkit.com and www.adhesives.org.

knowledge can be expressed, and they generally only contain a low number of adhesives or substrates. Finally, forward and backward chaining are not capable of providing all the functionality the expert needs. For instance, a situation might arise in which an adhesive is already pre-defined (e.g. left-over from a previous gluing operation), and the selection of a second substrate is required. While this selection requires the same knowledge, the expert tools are not capable of performing this operation.

## 3 Knowledge base paradigm

The core idea in the Knowledge Base Paradigm (KBP) (Denecker and Vennekens 2008) is to represent knowledge in a purely declarative way, independently of how it is to be applied. Knowledge is stored in a Knowledge Base (KB) and can be put to use via a multitude of inference tasks. In this way, the approach stimulates *knowledge reuse*, as multiple inference tasks can be used to solve multiple problems with the same knowledge.

### 3.1 IDP

The IDP (Imperative Declarative Programming) system (De Cat et al. 2018) is an implementation of the KBP. The knowledge in the KB is represented in a rich extension of First Order Logic (FOL), called FO($\cdot$) (pronounced *FO-dot*). It extends FOL with types, aggregates, inductive definitions and more. FO($\cdot$) is an expressive and flexible knowledge representation language, capable of modeling complex domains. The knowledge in a KB is structured in three kinds of blocks: *vocabularies*, *structures* and *theories*.

A *vocabulary* specifies a set of symbols. A symbol is either a type, predicate, or a function. A type is either a standard type such as the set of real numbers $\mathbb{R}$, or the name of an application-specific type, such as *Adhesive*. A predicate symbol either expresses a Boolean or a *relation* on one or more types, such as *BondSealing()* or *Available(Adhesive)*. A function symbol represents a function from the Cartesian product $T_1 \times \ldots \times T_n$ of a number of types to a type $T_{n+1}$. For example, the function *BondStrength* : *Adhesive* $\rightarrow \mathbb{R}$ maps each adhesive on its bond strength.

A *(partial) structure* specifies an interpretation for some of the symbols of a given vocabulary. A structure is *total* if it specifies an interpretation for each symbol of the vocabulary.

A *theory* contains a set of logical formulae in FO($\cdot$).

By itself, the KB cannot be executed: it is merely a "bag of knowledge," without information on how it should be used. The latest version of the IDP system, IDP-Z3 (Carbonnelle et al. 2022), supports many different inference tasks that can be applied to this knowledge. We will briefly go over the inference tasks that are relevant: **propagation**, **model expansion**, **optimization** and **explanation**. Given a partial interpretation $\mathcal{I}$ for the vocabulary of a theory $T$, **propagation** derives the consequences of $\mathcal{I}$ according to $T$, resulting in a more precise partial interpretation $\mathcal{I}'$. **Model expansion** extends a partial structure $\mathcal{I}$ to a complete interpretation $I$ that satisfies the theory T ($I \models T$). **Optimization** is similar to model expansion, but looks for the model with the lowest/highest value for a given term. Finally, **explanation** will, given a structure $\mathcal{I}$

which does not satisfy the theory T ($\mathcal{I} \not\models T$), find minimal subsets of the interpretations in $\mathcal{I}$ which together explain why $\mathcal{I}$ does not satisfy the theory.

### 3.2 Interactive consultant

The Interactive Consultant (Carbonnelle et al. 2019) is a graphical user interface for IDP-Z3, which aims at facilitating interaction between a user and the system. It is a generic interface, in the sense that it is capable of generating a view for any syntactically correct KB. In short, each symbol of the KB is represented using a *symbol tile*, which allows users to set or inspect that symbol's value. In this way, the GUI represents a partial structure to which a user can add and remove values. Each time a value is added, removed or modified, IDP's propagation is performed and the interface is updated: symbols for which the value was propagated are updated accordingly, and for the other symbols the values that are no longer possible are removed. In this way, a user is *guided* towards a correct solution: they cannot enter a value that would make the partial structure represented by the current state of the GUI inconsistent with the theory.

At any point in time the user can ask for an explanation of a value that was derived by the system, for example, when the user does not understand it or agree with it. The system will then respond with the relevant formulas and user-made assignments that lead to the derived value. In this sense, the tool is explainable, leading to more trust in the system.

A similar functionality is in place for the rare cases in which a user manages to reach an *inconsistent* state, that is, a set of assignments that can no longer be extended to a solution. While the IC removes values that have become impossible, it cannot do so for variables belonging to unbounded integer or real domains. For example, it is possible to input "Min Temperature = 20" followed by "Max Temperature = 10", as both of these variables can have any value between $-\infty$ to $+\infty$. In the case of an inconsistency, the interface alerts the user and explains why no solutions are possible by showing the relevant design choices and laws.

The Interactive Consultant interface has already successfully been used in multiple applications in the past (Aerts et al. 2022; Deryck et al. 2019).

### 4 Adhesive selector tool

This section outlines the creation and usage of the tool, and the main challenges that were faced in that process.

### 4.1 Knowledge acquisition

The creation of the knowledge base is an important element in the development process of knowledge-based tools. It requires performing knowledge acquisition, which is traditionally the most difficult step, as the knowledge about the problem domain needs to be extracted from the domain expert to be formalized by the knowledge engineer. While knowledge acquisition comes in many shapes and forms, we applied the *Knowledge Articulation* method (Deryck and Vennekens 2022). The central principle of this method
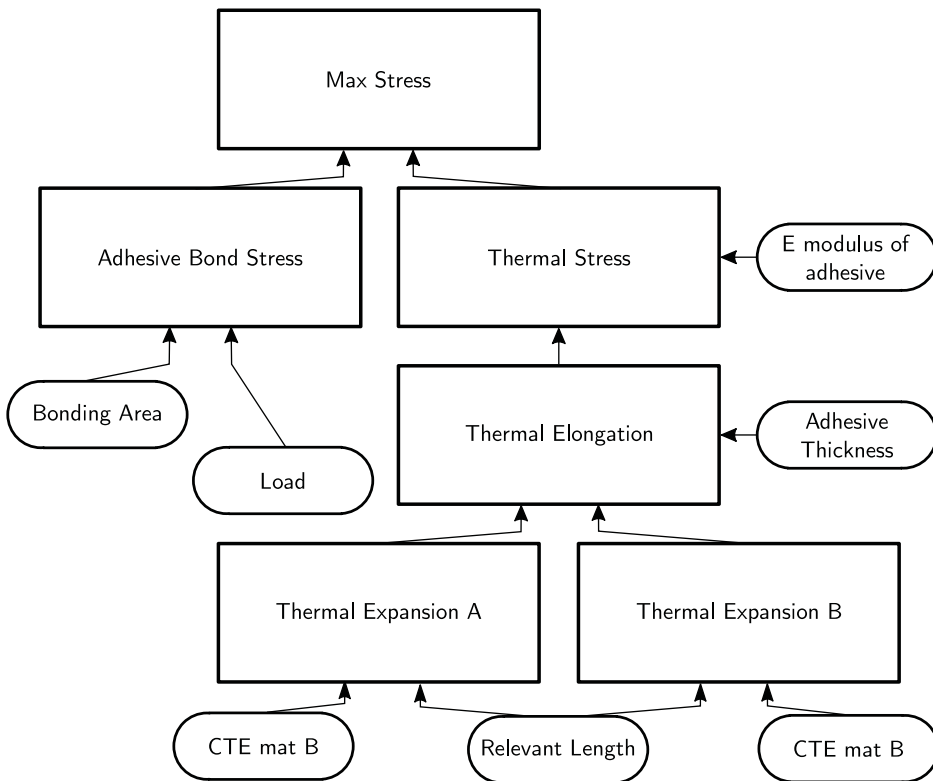
Fig. 1. Snippet of created DRD.

is to formalize knowledge in a common notation for both domain expert and knowledge engineer, so that both sides actively participate in the formalization process.

We started by organizing three knowledge articulation workshops, each lasting between three and 4 h. Each of these workshops was held with a group of domain experts. While typically a single domain expert would suffice for knowledge extraction, having a group present can help as an initial form of knowledge validation, as the experts discuss their personal way of working amongst themselves, before coming to a consensus. For the common notation we used Constraint Decision Model and Notation (cDMN) (Vandevelde et al. 2021), an extension of the Decision Model and Notation (DMN) standard (Object Modelling Group 2021). DMN is a user-friendly, intuitive notation for simple decision logic. Its main component is the *decision table*, which allows to define the value of a number of "output variables" in terms of a number of "input variables." Decision tables are structured together in a *Decision Requirements Diagram* (DRD), which is a graph that provides an overview of the total decision model by showing the connections between input variables (ovals) and decisions (rectangles). cDMN aims to increase the expressiveness of DMN (e.g. by adding constraints and quantification) while maintaining this user-friendliness.

The first workshop consisted of identifying all relevant adhesive selection parameters and using them to create an initial DRD, of which a fragment is shown in Figure 1. It is structured in a bottom-to-top way, similar to how the experts would reason: they start

(a)

| MinElongation | | |
|---|---|---|
| U | Support | MinElongation |
| 1 | free | $0.5 \times$ deltaLength / BondThickness |
| 2 | fixed | deltaLength / BondThickness |

Calculation of MinElongation

(b)

| BondStrength constraint | | |
|---|---|---|
| E* | Strength is Known | Max Stress |
| 1 | Yes | $\geq$ MinStrength |

BondStrength constraint table.

Fig. 2. Example cDMN tables.

by calculating the thermal expansions, and then work their way up to the calculation of the maximum stress.

During subsequent workshops, the rest of the model was fleshed out. This consists of *decision tables* and *constraint tables*. An example of such a decision table can be found in Figure 2a. In such a table, the "inputs" (in green, left) define the "outputs" (in light blue, right). Each row represents a decision rule, which *fires* if the values of the input variables match the values listed in the row. If a row fires, the value of the output is set accordingly. For example, if *Support = fixed*, then the *MinElongation* is calculated as *deltaLength/BondThickness*. The (U)nique hit policy of this table, indicated in the top left, means that the different rows must be mutually exclusive.

Figure 2b shows a constraint table (denoted by the *E\** in the top-left corner). In such a table, the output specifies a constraint that must hold if the input is satisfied. In other words, this table states that *if* the bond strength is known, *then Max Stress* should be higher than a minimum value. This differs from decision tables in that it does not define a specific value, but rather constrains its possible values.

After these three initial workshops, the cDMN model was converted into an FO(·) KB using the cDMN conversion tool. Since then, multiple one-on-one workshops were held between the knowledge engineer (first author) and the primary domain expert (second author) to further fine-tune the KB. Among other things, this included adding a list of adhesives, substrates, and their relevant parameter values, and further validating the knowledge. In total, the current version of the KB contains information on 55 adhesives and 31 substrates. For the adhesives, the KB contains 21 adhesive parameters, such as temperature resistances, strength and maximum elongation. Similarly, it contains 11 parameters for the substrates, such as their water absorption and their solvent resistance. These parameters are a mix of discrete and continuous: in total, 15 are continuous, and 17 are discrete.

### 4.2 Unknown adhesive parameters

One of the main challenges in the formalization of the KB was handling unknown adhesive data. Indeed, often an adhesive's data sheet does not list all of its properties. This raises the question of how the tool should deal with this: should the adhesive be excluded, or should it simply ignore the constraints that mention unknown properties? Together with

the experts we agreed on a third approach, in which we first look at the adhesive's family. Each adhesive belongs to one of 18 families, for which often some indicative parameter values are known. Whenever an adhesive's parameter is unknown, we use its family's value as an approximation. If the family's value is also unknown, then the constraint is ignored. This best corresponds to how the experts typically work.

This way of reasoning is formalized in the KB. For example, the constraint that an adhesive should have a minimum required bonding strength is written as follows:

$$\forall p \in param : Known(p) \Leftrightarrow (KnownAdhesive(p) \lor KnownFamily(p))$$

$$KnownAdhesive(strength) \Rightarrow BondStrength = StrengthAdhesive(Adhesive).$$
$$\neg KnownAdhesive(strength) \Rightarrow BondStrength = StrengthFamily(Family(Adhesive)).$$
$$Known(strength) \Rightarrow BondStrength \geq MinBondStrength.$$

with *StrengthAdhesive* and *StrengthFamily* representing respectively the specific adhesive's and its family's bonding strength. This approach is used for all 21 adhesive parameters.

One caveat to this approach is that IDP-Z3 currently does not support partial functions, that is, all functions must be totally defined. To overcome this, we assign the value $-1000$ to unknown parameter values, and define that the value is only *known* if it is different from this number. We chose $-1000$ as there is no adhesive parameter for which it is a realistic value.

$$\forall p \in param : KnownAdh(p) \Leftrightarrow StrengthAdhesive(Adhesive) \neq -1000.$$
$$\forall p \in param : KnownFam(p) \Leftrightarrow StrengthFamily(Adhesive) \neq -1000.$$
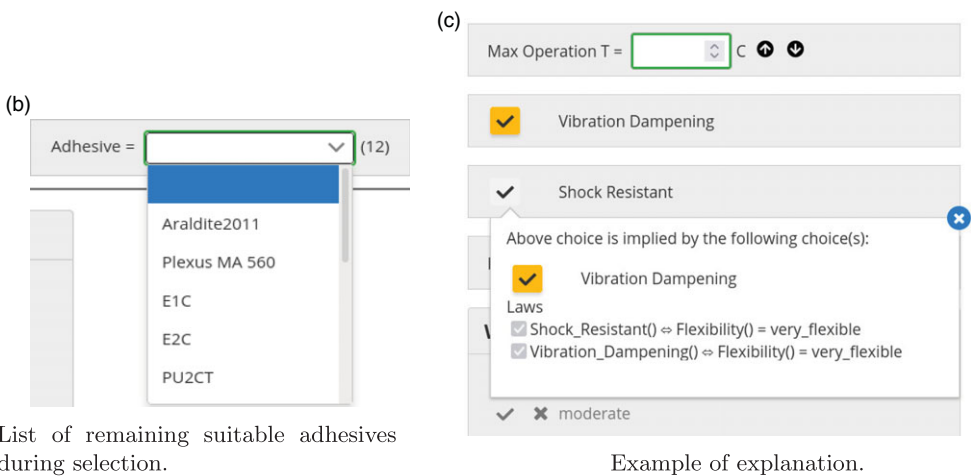
### 4.3 Interface

A crucial requirement of this application is the ability to interactively explore the search space. To this end, our tool integrates the Interactive Consultant to facilitate interaction with the KB. This interface makes use of several functionalities of the IDP system to make interactive exploration possible: the **propagation** inference algorithm is used to show the consequences of each choice, the **explain** inference is used to help the user understand why certain propagations were made, the **optimize** inference is used to compute the best adhesive that matches all of the choices made so far.

When using the interface, the user fills in symbol tiles, each representing a different symbol of the KB, and the system each time computes the consequences. For example, Figure 3a shows a segment of the interface in which a user set a maximum application temperature of 38°C as a requirement. To make it easier to navigate the symbol tiles, they are all divided in five *categories*: Performance, Production, Bond, Substrate A and Substrate B. In the top-right of the interface, the number of adhesives that remain feasible is shown: for example, after setting the temperature constraint, that drops from 55 to 12, as shown in Figure 3b.

The tool is also capable of generating two types of explanations. Firstly, if the user does not understand why a certain value was propagated, they can click on that value to receive a clarification, as demonstrated in Figure 3c. Secondly, if the user manages to reach an inconsistent state, the tool will try to help resolving the issue by listing what is

(a)
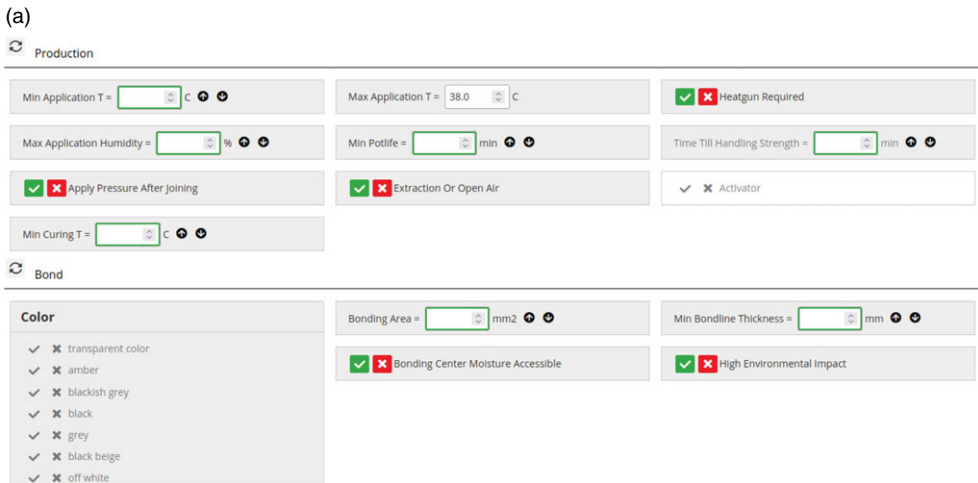
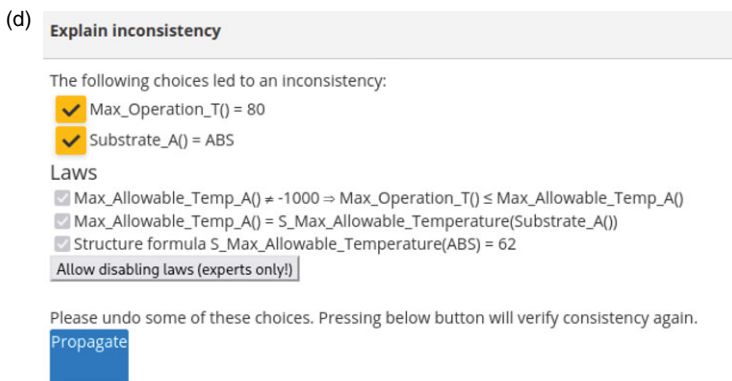

Some of the symbol tiles available in the interface, divided into two categories.

(c)



(b)



List of remaining suitable adhesives during selection.

Example of explanation.

(d)



Inconsistency window.

Fig. 3. Screenshots of the interface.

causing it. For example, Figure 3d shows an inconsistency in which a substrate is selected that cannot handle the required operating temperature.

Besides generating a list of all the adhesives that meet certain requirements, the tool can also find the optimal adhesive according to a specific criterion, such as lowest price or highest strength.

## 5 Preliminary validation

Initially, we performed three types of validation for this tool: a benchmark to measure the efficiency, a survey to measure the opinion of the adhesive experts and a discussion with the Flanders Make AI project lead.

*Benchmark.* In an initial benchmark, an adhesive expert was tasked with finding a suitable adhesive for an industrial use case which the company received. In total, it took the expert about 3 h to find such an adhesive, after delving through multiple data sheets. We then used our tool for the same use case, and were able to find the exact same adhesive within 3 min. Interestingly, the reasoning of the tool closely mimicked that of the expert: for example, they both excluded specific families for the same reasons.

*Survey.* After a demonstration of the tool to four adhesive experts, we asked them to fill out a short quantitative survey to better gauge their opinion. Their answers can be summarized as follows.

- The experts find the tool most useful for finding an initial list of adhesives to start performance testing with.
- The tool will be most useful for newer, less knowledgeable members of the lab. They can use the tool to learn more about the specifics of adhesive selection, or to verify if their result is correct.
- However, it is also useful for senior experts as they can discover adhesives which they have not yet used before.

The main criticism of the tool given by the experts is that more adhesives should be added, to make the selection more complete.

*Project lead discussion.* As part of a discussion with Flanders Make's project lead, who oversees multiple AI-related projects, they outlined their perception of our tool. They see many advantages. Firstly, as there is not much data available on the process of adhesive selection (e.g. previous use cases and the selected adhesives), and data generation is quite expensive, data-based approaches are not feasible. Therefore, building a tool based on a formalization of the knowledge they already have is very interesting. Secondly, by "storing" the expert knowledge formally in a KB they can retain this information, even when experts leave the company. Thirdly, having a formal representation also makes the selection process more uniform across different experts, who typically use different heuristics or rules-of-thumb. Lastly, they indicated that there is trust in the system, because the knowledge it contains is tangible. This makes it more likely that the experts will agree with the outcome of the tool.

The project lead also expressed that there is potential to maintain and extend this tool themselves, which would be a significant advantage compared to their other AI systems. However, we currently have not yet focused on this aspect.

## 6 User study

On top of the preliminary validation presented in Section 5, we performed a user study with members of the Flanders Make JML group. This study distinguishes itself from the preliminary validation in two ways: (1) all experts actually got to use the tool themselves and (2) it is a qualitative study instead of quantitative w.r.t. the input of the experts. Our main motivation for this study is to thoroughly validate the Adhesive Selector, with the following goals in mind:

1. Gauge the tool's effectiveness in a real-life setting.
2. Observe how new users interact with the tool.
3. Get feedback on the different aspects of the tool (interactivity, interface, explanations, ...).

For this validation, we asked the FM JML members to work out two real-life use cases, after which we performed one-on-one semi-structured interviews (SSI). In the following subsections we first elaborate on our methodology, then discuss the results of our interviews followed by describing the limitations of our approach.

### 6.1 Methodology

In our study, we held interviews with four members of the FM JML, who each possess knowledge on adhesives but have varying degrees of involvement in adhesive selection. Two interviewees are adhesive experts that often perform adhesive selection. The other two do not perform adhesive selection as part of their job, but are knowledgeable on glues in general. We included such "non-experts" to explore whether the tool is capable of making adhesive selection more accessible. We held an online one-on-one session with each of the interviewees. First the interviewee was asked to perform adhesive selection using the tool, and then we conducted a semi-structured interview to gather their opinions.

*Adhesive selection.* Two real-life bonding cases were selected by the second author to be used as test cases for the study: both are cases that the JML lab has received from companies in the past. In the first one, a plastic door needs to be glued to the body of an industrial harvester. Originally, the process of finding the correct glue took weeks, as the requirements are fairly tight. Moreover, the specification contains an inconsistency in which a higher temperature is required than allowed by the substrate, which is the same inconsistency as shown in Figure 3d. The second case details a join between a plastic component and an aluminum body, and is less challenging.

Both cases were presented to the participants as a short description of the gluing operation, together with a table containing the actual requirements. For the first example, the table lists requirements such as "Material A = Virgin ABS," "Gap filling of min 1 mm," "Application between 15°C and 35°C," etc. We have taken care here to specify require-

ments in the same terminology as the tool. In total, the use cases consist of nine and seven requirements respectively.

Before letting the experts work on the cases, we also gave a brief introduction on how to use the tool. We explained how to enter information, where they could find the list of possible adhesives and how to see that the tool was performing calculations.

We observed each tester during the selection process, to gather information on how they interacted with the tool. To better understand their thought process, the interviewees were encouraged to talk to themselves about their process as a way to elicitate their thoughts. We did not intervene while they were working out these use cases, even when errors were made, and only made suggestions on what to do if they got stuck or an unexpected bug popped up.

*Interviews.* Right after the selection, we held an interview with each participant. Because our goal is to explore the opinions of the JML members, we opted for a semi-structured interview set-up, for which we prepared four main questions to serve as a general guide:

1. Do you see a role for the tool in your job?
2. Do you feel that you understand what the tool does when you use it?
3. What was your experience working with the tool?
4. How would you compare our tool to the ones that you are used to working with?

In addition to these main questions, we asked additional questions to zoom in on specific aspects of the answers given by the participants. The interviews were led by the first author, with additional support from the second author. They were audio-recorded and transcribed, so that they could be analyzed thoroughly.

After transcribing the recordings, we followed the guidelines of Richards and Hemphill (2018) and performed *open coding* followed by *axial coding* (Corbin and Strauss 1990). Here, the goal is to identify various *codes* that pop up during the interviews, and then further group them into several main categories. These two steps were also performed separately by an external member of the research group, after which the results were compared and adapted to mitigate biases (*consensus coding*).

### 6.2 Results

Based on the interview transcripts, we identified 26 codes in total. Table 1 shows an overview of the interview statistics. The last column of this table shows the *Code distribution*, calculated as the cumulative percentage of codes discovered after each interview. This is an important parameter that indicates whether *data saturation* is reached, that is, the point at which the same themes keep recurring and new interviews would not yield new results. According to Guest et al. (2020), this point is reached when the difference in code distribution between the current and previous interview is $\leq 5\%$ (i.e. less than 5% of new information was found). As the information threshold between the third and fourth interview is 4%, we conclude that we have reached data saturation. A more detailed table showing the codes per interview is included in the Appendix.

We sub-divided the codes into five main themes: Knowledge, Interactivity, Expert, Interface and Explainability. These themes, together with their codes, are visualized in the graph in Figure 4. We will now briefly go over each theme and highlight the most important findings.

Table 1. *Interview statistics*

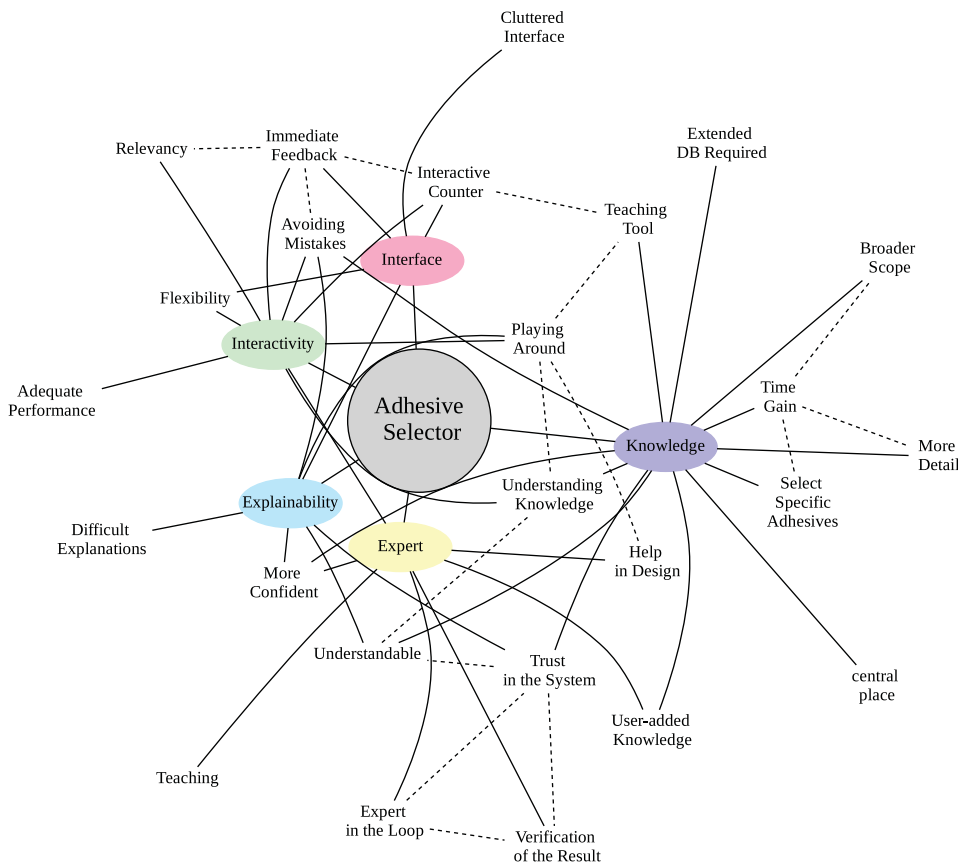| Use case duration (min) | Interview duration (min) | Words | total codes | Code distribution |
|---|---|---|---|---|
| 31 | 26 | 2092 | 15 | 58% |
| 55 | 39 | 2107 | 16 | 88% |
| 66 | 47 | 3497 | 19 | 96% |
| 41 | 49 | 3680 | 13 | 100% |



Fig. 4. Graph showing connections between the interview themes and their codes.

### 6.2.1 Knowledge

A major advantage of our tool is that it reasons on the knowledge of the adhesive experts "like they themselves would." This is affirmed by the participants, who particularly liked its level of detail in two ways. Firstly, instead of being limited to selecting adhesive families, as is the case in other selector systems, our tool can help them find specific adhesives. Secondly, the number of parameters available in the tool is unparalleled, allowing for a more fine-grained search.

> "[...] the good thing that I like about this tool is that it's quite detailed. It's really one or multiple steps further, multiple steps deeper than these other tools. [...] So there it's super helpful, definitely."

Similarly, the tool can also efficiently reason on a broader number of specific adhesives than an expert. As the latter typically knows some approximate parameter values by heart for a handful of glues and families, they tend to look at these adhesives first before widening their search to others when needed. Here, the Adhesive Selector can help them to find additional adhesives that they would not have considered without the tool, while also saving them from having to manually go through their data sheets.

> "You might encounter situations in which you find an adhesive that you hadn't thought about before. That quite increases your search scope, I think. It could also help in mitigating bias."

> "If you want to understand the glue that you want to use you have to read a lot of datasheets, while having this tool I think optimized the time in a way that is crazy."

The main criticism expressed by the participants is that the number of specific adhesives in the knowledge base is still rather limited. While 55 adhesives is a good start, adding more adhesives will definitely make the tool more effective. Similarly, some participants remarked that the knowledge base should be with more expert knowledge such as the environmental impact and re-usability of adhesives.

### 6.2.2 Interactivity

Throughout the interviews, interactivity was an often recurring topic. For example, one aspect that the participants all appreciated was that they could immediately see the effects of entering requirements in the interface.

> "Interactively choosing the glue now feels like online shopping – I can select more options, and see the total number of adhesives go down until I have entered all requirements."

Besides making the number of adhesives go down, each time a new requirement is entered the interface also grays out parameters that have become irrelevant, and removes parameter values that are no longer possible. This immediate feedback helps preventing mistakes in the selection.

> "I think it's easier to spot if there are some problems, like the ones that popped up."

> "It's safer, you avoid wasting materials and time."

Moreover, the immediate feedback also allows the experts to "play around" with the knowledge in the tool. In this way, they can get a feel for the effect of certain parameter values on the suitable adhesives.

> "Us engineers typically want to play with things. They want to see what happens when they change something, thereby implicitly performing a sensitivity analysis."

However, some participants felt that always having to enter the requirements one-by-one is too inefficient. While they all agreed that the tool is sufficiently fast, they stated that they would like to "bulk update" choices to be more efficient in cases where they already know that they are correct.

### 6.2.3 Expert

As our tool is designed to support JML's experts based on their own knowledge, they have played a big role in the creation of the tool and will continue to play a big role in its further existence. It is necessary to always have an expert in the loop: the tool cannot be used by laymen, who lack the specific knowledge required to extract requirements from a description of a use case and who cannot understand the technical jargon.

Besides supporting the experts in making suitable selections, the participants described two other ways in which the Adhesive Selector can help them. Firstly, the tool covers sufficient knowledge to also assist in the design of the entire joint. This is different from adhesive selection in that many "environmental parameters" are still left open, for example, substrates might need to be picked, a decision on joining method needs to be made, etc.

> "This tool can already help me to list all these requirements [required for the dimensioning of the joint] so that I have a bird's eye view of the whole design."

The second additional use of the tool identified by the participants is its potential to be used as a teaching tool. For instance, newer members of the lab could use it to gain their footing when starting out.

> "You can use it to teach people 'If you select these requirements, these are the consequences, which means you can no longer use these adhesives'."

> "[The tool] can also give some confidence, if they say 'I would select this' and the tool confirms it, you would feel more certain in your selection"

## 6.2.4 Interface

Having the participants talk out loud while working the use cases proved to be a valuable source of information, allowing us to gain insight in how someone without prior experience interacts with the tool. Some of these insights are fairly minor, for example, that we should order the values of the drop down lists alphabetically and support folding in/out entire categories to make the interface easier to navigate. Other insights are more major, such as how much difficulty the participants experienced with the structure of the interface – due to the large number of parameter tiles (78 in total), it typically took them around 10 s to find the right one. Moreover, the participants often lost overview of which choices they had already entered in the system, as these are spread all over the interface.

One positive note that the experts really liked was the automatically updating counter showing the remaining number of suitable adhesives:

> "I can look at the number of adhesives and see that we are converging, converging, converging. So it's fun, to say it like that."

This ties in nicely with the idea of the Adhesive Selector as a didactical tool: if a choice rules out many adhesives at once, the user can assume it is more "important" than a choice that only removes a handful. Moreover, another participant suggested using diagrams to annotate symbols with their intended meaning, for example, a diagram showing two bonded substrates with arrows pointing to the "bond line," the "bonding surface," etc. This would make the tool more self-explanatory to people less familiar with these terms, such as the newer members of the team.

## 6.2.5 Explainability

While explainability is one of the focuses of our knowledge-based approach, the experts were not yet fully convinced of this functionality. When prompted with the "inconsistency" window (Figure 3d), none of the participants knew what to do. Only two of them quickly understood what the cause of the inconsistency was, but none were able to resolve it by themselves. As one participant later remarked in an interview:

> "There were multiple sentences below each other, I didn't know if it was three remarks or a single one. [...] I was confused, and could not see the information I needed"

In other words, they had some difficulties navigating the inconsistency window: partly due to its layout, but also due to the complexity of the knowledge. However, they did appreciate the potential that this feature holds, for example to assist in experimenting with the knowledge.

The difficult explanations also did not have an impact on the expert's trust in the system. Indeed, it seems that it is more important to know that its behavior is derived from the lab's own knowledge, than to actually 100% understand the explanations.

> "I trust it, because it contains our knowledge. So, well, I do trust it, but only because I know it's ours."

*Other results.* A participant with a slightly more business background pointed out an interesting result that could not be codified under the other themes. They see the tool as a "uniform way of collecting data relevant to adhesive selection," which could help drive the team's decision making. Examples of relevant information are the use cases performed by experts, the substrates that were used, which adhesives are typically picked, what constraints were present, etc. This data could be useful for:

- finding trainings specifically for the most commonly picked adhesives
- identifying target industries that could also be helped by JML
- deciding which adhesives to keep in stock, and what equipment should be purchased

This is a unique take that we had not yet considered for our tool.

### 6.3 Limitations of the study

As with any study, ours is not without its limitations. The main limitation of this study is our low number of interviewees. However, as pointed out in Section 6.2, we do reach data saturation when looking at the code distribution of our interviews. Therefore, we feel that the low number of interviewees does not have a major impact.

Another limitation is that all interviewees are from the same organization and would therefore have a "common" approach to selecting glues, while external experts might have a different focus that could result in additional feedback. However, as the tool has specifically been developed for use within Flanders Make JML, we believe that this feedback would be less relevant.

During the testing of the tool, we continually observed the interviewees to study their interaction with the tool. Because of this, the users might have felt pressured to work in a more "efficient" way, as a form of the Hawthorne effect (McCarney et al. 2007). In fact, one participant said so explicitly, when talking about their experience with the inconsistency window:

> "I think that if I were alone, in the lab, I would have taken more time to read the pop-up. I wanted to be a bit quick."

Because we were observing, we might have inhibited the users to truly "play around" with the tool and test it to their heart's content.

### 7 Lessons learned

Typically, knowledge acquisition is a time-consuming and difficult process. We have found that the use of cDMN as a common notation can help facilitate this process. The use of a formal representation that the experts can also understand helps to keep them in the loop and allows them to actively participate in the formalization process. This way of working is less error-prone, as it functions as a preliminary validation of the knowledge.

After our three initial workshops, we mainly held one-on-one meetings with one of the experts to add information on the adhesives, and to further fine-tune the knowledge. This resulted in a tight feedback loop, which turned out to be a key element in our formalization. Indeed, thanks to thorough examinations of the tool by the expert, we were able to discover additional bugs in our KB. Here, the Interactive Consultant was of paramount importance: each time the KB was modified, the expert could immediately play around with it using the generic interface. In this way, the knowledge validation of the tool could happen immediately after the modifications, allowing for a swifter detection of any errors.

Having knowledge in a declarative format, independent of how it will be used, has multiple advantages. To begin with, it allows using the knowledge for multiple purposes, even when this initially might not seem useful. Furthermore, it increases the experts' trust in the system, as it reasons on the same knowledge as they do, and is interpretable.

The main advantage of using IDP-Z3 does not lie in any one of its inference algorithms, but rather in the fact that it allows all of the functionalities that are required for interactive exploration of the search space to be performed by applying different inference algorithms to a single KB.

The validation of the tool by the actual end-users proved to be a source of valuable feedback. Through our observations, we have gained insights on how the users interact with the tool. By means of semi-structured interviews, we gathered their opinions, thoughts and suggestions. These two sources of input combined will help shape further development of our tool.

## 8 Conclusions and future work

This paper presents the Adhesive Selector, a tool to support adhesive selection using a knowledge-based approach. The Knowledge Base was constructed by conducting several workshops and one-on-one meetings, using a combination of DMN and cDMN. Our current iteration of the tool contains sufficient knowledge to assist an expert in finding an initial list of adhesives. Compared to the state of the art, it is declarative, more explainable, and more extensive. The KB is also not limited to just adhesive selection, but can also be used to perform other related tasks.

In future work, we plan on converting the entire FO($\cdot$) KB into cDMN, and evaluating its readability and maintainability from the perspective of the domain experts. Besides this, we intend to test the tool using more real-life use cases, to quantify the gain in efficiency. Additionally, we are also collaborating with an external research group to develop an AI-based tool capable of extracting adhesive parameter values from data sheets, to efficiently add more adhesives to our KB. Flanders Make is planning on using the tool in production soon, and expanding the list of available adhesives is the only remaining bottle-neck. Note that such information is only a part of the knowledge that is required for selecting a suitable adhesive, and other forms of expert knowledge will continue to be added throughout the lifetime of the tool through the methods described in the text.

## Competing interests

## Acknowledgements

## References

ALLEN, R. B. AND VANDERVELDT, H. 1995. Expert-system selects adhesives for composite-material joints. *Adhesives Age 38*, 7, 16–19.

AERTS, B., DERYCK, M. AND VENNEKENS, J. 2022. Knowledge-based decision support for machine component design: A case study. *Expert Systems with Applications 187*, 115869.

CARBONNELLE, P., AERTS, B., DERYCK, M., VENNEKENS, J. AND DENECKER, M. 2019. An Interactive Consultant. In *Proceedings of the 31st Benelux Conference on Artificial Intelligence (BNAIC 2019) and the 28th Belgian Dutch Conference on Machine Learning (BENELEARN 2019), Brussels, Belgium, November 6–8, 2019*, vol. 2491. *CEUR Workshop Proceedings*. CEUR-WS.org.

CARBONNELLE, P., VANDEVELDE, S., VENNEKENS, J. AND DENECKER, M. 2022. IDP-Z3: a reasoning engine for FO(.). In *FOSDEM 2023*.

CORBIN, J. M. AND STRAUSS, A. 1990. Grounded theory research: Procedures, canons, and evaluative criteria. *Qualitative Sociology 13*, 1, 3–21.

DA SILVA, L. F. M., CHSNER, A. AND ADAMS, R. D. 2018. *Handbook of Adhesion Technology*, vol. 1. Springer.

DE CAT, B., BOGAERTS, B., BRUYNOOGHE, M., JANSSENS, G. AND DENECKER, M. 2018. Predicate logic as a modeling language: The IDP system. In *Declarative Logic Programming: Theory, Systems, and Applications*, M. Kifer, and Y. A. Liu, Eds. ACM, 279–323.

DENECKER, M. AND VENNEKENS, J. 2008. Building a knowledge base system for an integration of logic programming and classical logic. In *Logic Programming*, M. Garcia de la Banda and E. Pontelli, Eds., vol. 5366. LNCS. Springer Berlin Heidelberg, Berlin, Heidelberg, 71–76.

DERYCK, M., DEVRIENDT, J., MARYNISSEN, S. AND VENNEKENS, J. 2019. Legislation in the knowledge base paradigm: interactive decision enactment for registration duties. In *Proceedings of the 13th IEEE Conference on Semantic Computing*. IEEE, 174–177.

DERYCK, M. AND VENNEKENS, J. 2022. An integrated method for knowledge management in product configuration projects. In *Towards Sustainable Customization: Bridging Smart Products and Manufacturing Systems*, A.-L. Andersen, R. Andersen, T. D. Brunoe, M. S. S. Larsen, K. Nielsen, A. Napoleone and S. Kjeldgaard, Eds. Cham: Springer International Publishing, 860–868.

EWEN, J. C. K. 2010. Key issues in selecting the right adhesive. In *Advances in Structural Adhesive Bonding*, D. A. DILLARD, Ed. Woodhead Publishing in Materials. Woodhead Publishing, 3–19.

GUEST, G., NAMEY, E. AND CHEN, M. 2020. A simple method to assess and report thematic saturation in qualitative research. *PLoS One 15*, 5, e0232076.

KANNAN, T. AND PRABU, S. S. 2004. Expert system for selection of adhesives. In *Proceedings of the Recent Developments in Materials Processing Conference*.

LAMMEL, C. AND DILGER, K. 2002. Software for a rule-based adhesive-selection system. *Adhesives & Sealants Industry 9*, 5, 42–43.

LEES, W. AND SELBY, P. 1993. The PAL program mark II. *International Journal of Adhesion and Adhesives 13*, 2, 120–125.

McCARNEY, R., WARNER, J., ILIFFE, S., VAN HASELEN, R., GRIFFIN, M. AND FISHER, P. 2007. The Hawthorne effect: A randomised, controlled trial. *BMC Medical Research Methodology 7*, 1, 1–8.

MEYLER, K. L. AND BRESCIA, J. A. 1993. Design of a computer expert system for adhesive selection using artificial intelligence techniques. Technical report, Army armament research development and engineering center Picatinny Arsenal.

MOSELEY, L. AND CARTWRIGHT, M. 1992. The development of an expert system for operational use in the selection of industrial adhesives. *Engineering Applications of Artificial Intelligence 5*, 4, 319–328.

OBJECT MODELLING GROUP. 2021. Decision model and notation v1.3. URL: `http://www.omg.org/spec/DMN/`

RICHARDS, K. A. R. AND HEMPHILL, M. A. 2018. A practical guide to collaborative qualitative data analysis. *Human Kinetics 37*.

SU, Y., SRIHARI, K. AND ADRIANCE, J. 1993. A knowledge update mechanism for an adhesive advisor. *Computers & Industrial Engineering 25*, 1–4, 111–114.

VANDEVELDE, S., AERTS, B. AND VENNEKENS, J. 2021. Tackling the DM challenges with cDMN: A tight integration of DMN and constraint reasoning. *Theory and Practice of Logic Programming 23*, 3, 535–558.

VANDEVELDE, S., JORDENS, J., VAN DONINCK, B., WITTERS, M. AND VENNEKENS, J. 2022. Knowledge-based support foradhesive selection. In *Logic Programming and Nonmonotonic Reasoning*, G. Gottlob, D. Inclezan and M. Maratea, Eds. Cham: Springer International Publishing, 445–455.

# Appendix

Table 2 shows an overview of each code and the interviews in which it appeared, and Table 3 contains the "code book" explaining each of the codes.

While we cannot release the knowledge base itself, as it contains sensitive data from our external partner, the tools and techniques used in this paper are open-source and available online. We refer anyone interested in trying out IDP-Z3[2], the Interactive Consultant[3] or cDMN[4] to their online resources.

---

[2] `https://www.IDP-Z3.be` and `https://gitlab.com/krr/IDP-Z3/`
[3] `https://interactive-consultant.idp-z3.be/`
[4] `https://www.cdmn.be` and `https://gitlab.com/EAVISE/cdmn/cdmn-solver`

Table 2. *Discovered codes per participant, and the data saturation*

| Code | 1 | 2 | 3 | 4 |
|---|---|---|---|---|
| Cluttered interface | ✓ | ✓ | ✓ | ✓ |
| Interactive counter | ✓ | ✓ | ✓ | |
| Time gain | ✓ | ✓ | ✓ | ✓ |
| Avoiding mistakes | ✓ | | | |
| Playing around | ✓ | | ✓ | |
| More confident | ✓ | | ✓ | ✓ |
| Understandable | ✓ | | | |
| Trust in the system | ✓ | ✓ | ✓ | ✓ |
| Verification of the result | ✓ | ✓ | ✓ | ✓ |
| Adequate performance | ✓ | ✓ | ✓ | ✓ |
| Broader scope | ✓ | | ✓ | ✓ |
| Understanding knowledge | ✓ | | | ✓ |
| Teaching tool | ✓ | | ✓ | ✓ |
| Explainability | ✓ | ✓ | ✓ | |
| Difficult explanations | ✓ | ✓ | ✓ | ✓ |
| Extended DB required | | ✓ | ✓ | |
| Select specific adhesives | | ✓ | ✓ | |
| Help in design | | ✓ | | |
| Expert in the loop | | ✓ | ✓ | ✓ |
| Interactivity | | ✓ | ✓ | |
| Relevancy | | ✓ | | |
| More detail | | ✓ | ✓ | ✓ |
| User-added knowledge | | ✓ | | |
| Flexibility | | | ✓ | |
| Immediate feedback | | | ✓ | |
| Central place for knowledge | | | | ✓ |
| Total codes | 15 | 16 | 19 | 13 |
| Data saturation | 58% | 88% | 96% | 100% |

Table 3. *A "code book" elaborating on the meaning of each code*

| Code | Explanation |
| --- | --- |
| Cluttered interface | The interface's structure seems chaotic. |
| Interactive counter | The counter denoting the number of suitable adhesives. |
| Time gain | Experts can more efficiently select an adhesive. |
| Avoiding mistakes | The tool prevents making mistakes. |
| Playing around | Interacting with the knowledge allows the expert to "play around" with it. |
| More confident | The experts are more sure of their selection using the tool. |
| Understandable | The interface presents information in a sufficiently understandable way |
| Trust in the system | The experts trust the choices of the system. |
| Verification of the result | The resulting short-list of adhesives will always need to be verified. |
| Adequate performance | The loading time of the tool is within reason. |
| Broader scope | The experts discover more adhesives. |
| Understanding knowledge | The tool helps understand the knowledge. |
| Teaching tool | The tool is well-suited for teaching newer members of the lab. |
| Explainability | Explainability is an important factor. |
| Difficult explanations | The automatically generated explanations are difficult. |
| Extended DB required | More adhesives should be added to the system. |
| Select specific adhesives | The AS allows selecting specific adhesives, whereas the other tools only select families. |
| Help in design | The tool can also play a role in joint design. |
| Expert in the loop | An expert is always required to stay in the loop. |
| Interactivity | Interactivity is a crucial aspect of the tool. |
| Relevancy | Detecting which parameters are still relevant aids in selecting adhesives. |
| More detail | The tool is more detailed than any other in terms of possible parameters. |
| User-added knowledge | The experts would like to add their own knowledge to the system. |
| Flexibility | The tool adapts well to the expert's way of working. |
| Immediate feedback | The experts appreciate that the interface updates automatically. |
| Central place for knowledge | All knowledge for adhesive selection is stored in one location. |