

# Supporting the Problem-Solving Loop: Designing Highly Interactive Optimisation Systems

Jie Liu, Tim Dwyer, Guido Tack, Samuel Gratzl, Kim Marriott

**Abstract**—Efficient optimisation algorithms have become important tools for finding high-quality solutions to hard, real-world problems such as production scheduling, timetabling, or vehicle routing. These algorithms are typically “black boxes” that work on mathematical models of the problem to solve. However, many problems are difficult to fully specify, and require a “human in the loop” who collaborates with the algorithm by refining the model and guiding the search to produce acceptable solutions. Recently, the Problem-Solving Loop was introduced as a high-level model of such interactive optimisation. Here, we present and evaluate nine recommendations for the design of interactive visualisation tools supporting the Problem-Solving Loop. They range from the choice of visual representation for solutions and constraints to the use of a solution gallery to support exploration of alternate solutions. We first examined the applicability of the recommendations by investigating how well they had been supported in previous interactive optimisation tools. We then evaluated the recommendations in the context of the vehicle routing problem with time windows (VRPTW). To do so we built a sophisticated interactive visual system for solving VRPTW that was informed by the recommendations. Ten participants then used this system to solve a variety of routing problems. We report on participant comments and interaction patterns with the tool. These showed the tool was regarded as highly usable and the results generally supported the usefulness of the underlying recommendations.

**Index Terms**—Interactive optimisation, Interface design, Usability, Interactive systems and tools, Vehicle routing

## 1 INTRODUCTION

Automatic optimisation is increasingly used to find high-quality solutions to hard, real-world resource allocation and scheduling problems, such as timetabling [55], logistics [6], scheduling of scientific observations [54], or planning of medical procedures [9]. The problem is first modelled mathematically and then a constrained optimisation solver is used to find the best—or at least a very good—solution to the problem [42].

The traditional view of the optimisation solver is that it is a “black box” which does not support user interaction. However, the optimisation community has found that in practice many real-world problems require the user to actively engage in the solution process. One of the main reasons for such *interactive optimisation* is that a mathematical model necessarily simplifies the real-world problem [2, 17]. Interaction allows the user to bring their additional knowledge into the solution process and means that the tool is less brittle and can more flexibly adjust to unforeseen situations. In addition, interaction builds the user’s trust and confidence in the solver [36].

The design of visual interfaces that support people to work effectively with optimisation systems can be seen as an application of Visual Analytics [35]. Here we present and provide an initial evaluation of the first design recommendations for such interfaces. Our specific contributions are fourfold.

1. Nine recommendations informed by the *problem-solving loop* [35]. This loop is a theoretical framework for understanding the high-level user goals and tasks in interactive optimisation and is an analogue of the sense-making loop widely used to understand visual analytics. The recommendations cover visualisation of solutions and constraints, interaction with the solver, as well as comparison of solutions and provenance of solutions.
2. A review of representative interactive optimisation tools to determine how well these recommendations are supported by current tools. While some recommendations were well supported, most—such as showing a gallery of solutions, side-by-side comparison of

solutions, showing solution provenance or feedback on the solution process—were not.

3. The design and implementation of an interactive optimisation tool that exemplifies all of the design recommendations. Our tool solves representative real-world problem with diverse constraints: the vehicle routing problem with time windows (VRPTW).
4. A qualitative user study with 10 participants evaluating the interactive tool for solving VRPTW. The study explored: the usability of the tool; its flexibility in adjusting to new scenarios; the tool’s support for our recommendations; and the usefulness of those recommendations as evidenced by the tool’s use.

This research informs the design of future optimisation tools. More broadly it informs the design of other kinds of AI-based decision-support tools, and visual analytics tools in general.

## 2 BACKGROUND

**Interactive Optimisation:** Optimisation uses computational techniques—such as integer programming, simulated annealing or constraint programming—to find the best (or at least a good) solution to a mathematical model of a decision problem. The model specifies *decision variables* for which the optimisation technique will find values, an *objective function* to measure the quality of a solution, and *constraints* that restrict the values that the decision variables can take. Usually, we create a general model which is instantiated to a particular problem by assigning values to the model’s *parameters*.

For example, in this paper we consider the Vehicle Routing Problem with Time Windows (VRPTW). In this problem the model’s parameters specify the number of vehicles available to visit customers, the time windows in which each customer must be visited, and the distance and travel time between customers and from each customer to the depot. A solution to the problem assigns each customer to a vehicle, and determines the order of customer visits for each vehicle. The constraints are that customers can only be visited in their time windows, each vehicle can only visit one customer at a time and must have time to travel between consecutive customers. The objective is to minimise the total distance travelled by all vehicles.

There is now a recognition by the optimisation community of the need in many applications for optimisation systems that are interactive rather than fully automatic “black boxes”, for which the user simply provides values for the model’s parameters and then waits for the system to compute a solution. A recent survey paper [38] clarifies the rationale for involving the user. For many real-world problems it is unrealistic to fully model the problem. Allowing the user to modify

• Jie Liu, Tim Dwyer, Samuel Gratzl, Kim Marriott and Guido Tack are with Monash University. E-mail: {jliu120@gmail.com, {tim.dwyer3|guido.tack|kim.marriott}@monash.edu, sam@sgratzl.com

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

constraints and adjust trade-offs in the objective function means that they can reduce the gap between the model and the real-world problem. It also provides flexibility as the user can modify the model to handle unexpected situations. Other benefits of interaction are that it may allow the user to guide the optimisation algorithm to find a better solution and can help the user build an appropriate level of trust in the system.

Most research into interactive optimisation has focused on algorithms or applications. For instance, the aforementioned survey [38] almost totally ignores visualisation and interaction techniques, the user experience or user evaluations. Two notable exceptions are an early review by Jones [28] of visualisation usage in optimisation and a survey by Miettinen [39] of visualisations used for multi-criteria optimisation. Other research, e.g. Goodwin et al. [24], has investigated how to use visualisation to diagnose and refine optimisation models. However, this is to support the initial development of the model by an expert in optimisation, not the subsequent use of the model by the end-user.

One area that has received some attention is whether interaction engenders trust in optimisation systems. Liu et al. [35] found that participants gained trust through the use of interactive optimisation tools from a small qualitative study in brachytherapy. And in two controlled studies Liu et al. [36] found that feedback about the solution process led to over-trust, but allowing users to manipulate solutions led to better-calibrated trust.

A number of user studies have evaluated the effectiveness of interactive optimisation systems [4, 8, 10, 12, 13, 15, 17, 18, 23, 25, 27, 30, 32, 37, 45, 50, 52, 56]. In particular, Caldas and Santos [13] developed a user-guided interactive system for daylighting design and found the proposed system could produce good-quality solutions. Similarly, Matejka et al. [37] presented Dream Lens, an interactive visual analysis tool to explore and visualise generative design. Another user study evaluated an interactive optimisation system developed by Coffey et al. [15] that proposed a user-guided approach to allow users to directly navigate the simulation model and change the design, while the system adjusted the underlying parameters correspondingly. Dayama et al. [18] described an interactive optimisation system for layout design in which the tool suggested alternative layouts in an example gallery. Klau et al. [32] proposed an interactive optimisation system based on the human-guided simple search (HuGSS) framework to solve a particular kind of vehicle routing problem and reported that better solutions could be found by using the system.

However, to the best of our knowledge there has been no attempt to develop or evaluate general recommendations for the design of visualisation and interactions in interactive optimisation systems.

**Visual Analytics:** Another essential tool for decision support is visual analytics. This uses automatic analysis empowered by visualisation and interaction techniques to make sense of data [31]. There are great similarities between visual analytics and interactive optimisation as both research fields aim to employ both humans and computers to solve complex real-world problems. Indeed, in a sense interactive optimisation is a visual analytics task in which the user is trying to understand the space of possible solutions to a decision problem [35]. Unlike interactive optimisation, usability and the creation of design guidelines has been an ongoing focus of visual analytics research.

There exist several models or frameworks describing high-level tasks and processes in visual analytics. Pirolli and Card [46] proposed a sense-making loop which captured the processes and steps employed by analysts when analysing and making sense of information. Keim et al. [31] presented the visual analytics process as a knowledge discovery framework describing typical interactions and tasks. Ceneda et al. [14] introduced a model, extended from Van Wijk's [59] generic visualisation model, to describe the opportunities of using automatic guidance to support user analysis in visual analytics. This model was further extended by Collins et al. [16] with a focus on improving the efficiency of the analytical process utilising different automatic guidance. Dudley et al. [21] gives a workflow for interactive machine learning (ML).

**Guidelines:** While currently there are no design guidelines for interactive optimisation systems, guidelines exist for other kinds of interactive applications. Endsley [22] gives general guidelines for the design of autonomous and semi-autonomous systems. She highlights the im-

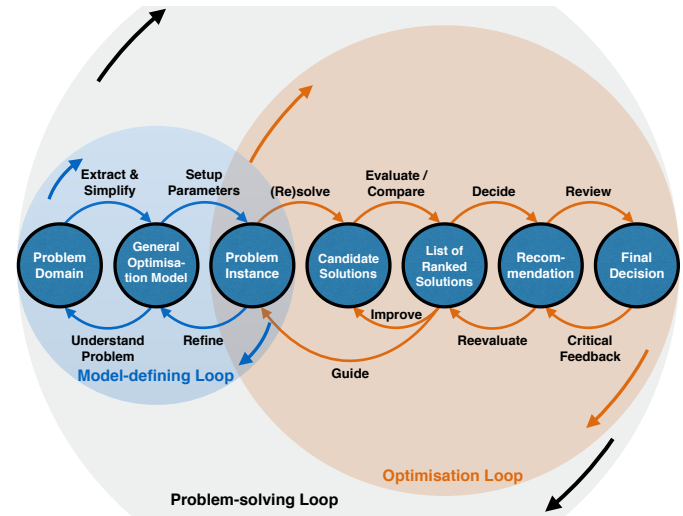


Fig. 1. The problem-solving loop. Image from [35].

portance of user involvement and control, recommending that tools support and assist decision making but do not make the final decision. Research into interactive ML highlights the need for meaningful explanations of system behaviour [34], the need to engage the user, provide effective data representations, and to exploit interactivity and provide rich interactions [21]. General design recommendations for AI systems emphasize the need for transparency in what the system can do and why it has made a particular decision or performed a particular action [1]. Nodalo et al. [43] also emphasizes transparency and the need to use simple, understandable visualisations. Mühlbacher et al. [40] presents different models of interactive visualisation with analytic software and strategies for increasing user involvement.

**Problem Solving Loop:** If we are to develop design recommendations for interactive optimisation tools it is crucial to understand the fundamental tasks and processes in interactive optimisation. Hillier [26] summarised the general steps in designing an automatic optimisation system, which they named the 'design process'. Shim [53] presented and explained the general decision-making process of decision support systems. However, both Hillier and Shim only considered fully automatic optimisation rather than interactive optimisation.

Recently, Liu et al. [35] proposed a theoretical framework called the *problem-solving loop*, see Fig. 1, articulating the user goals and tasks in interactive optimisation. It was based on feedback from health professionals on an interactive optimisation application in treatment planning for prostate cancer. The problem-solving loop has two main loops: the *model-defining loop* and the *optimisation loop*. The model-defining loop captures the process of creating an abstract mathematical model of the general problem, while the optimisation loop captures the interactive use of the model to solve a particular instance of the problem. However, Liu et al. [35] did not evaluate this framework or use it as the basis for developing design recommendations for interactive optimisation systems. This is the focus of this paper. We now examine the loop more closely.

**Model Defining Loop:** There are four tasks in the model-defining loop. The two leftmost tasks *Understand Problem* and *Extract & Simplify* capture the initial creation of the generic model by the (usually) expert optimisation modeller. These tasks do not need to be supported by an interactive optimisation tool as the assumption is that they were completed when the tool was first created. The other two tasks, however, should be supported in an interactive user interface. They are:

- *Setup Parameters:* Determining the parameter values for the instance of the problem to be solved;
- *Refine:* Adding/removing/changing constraints and/or objectives in the optimisation model.

**Optimisation Loop:** The optimisation loop has eight tasks. The right-most four tasks do not need to be directly supported by the interactive optimisation tool as they are concerned with deciding on the best solu-

tion out of a ranked list of solutions and then the review of the solution by other stakeholders. The remaining four tasks, however, should be supported by the tool:

- *(Re)solve*: Using the constraint solver to solve the generic model with problem-specific parameters to generate a pool of candidate solutions;
- *Evaluate/Compare*: Showing a specific single solution and its associated constraint(s) and objective(s) to support the solution evaluation process; Or comparing multiple candidate solutions to give a ranked list of these solutions to better understand their trade-offs;
- *Improve*: Manipulating a solution to improve it;
- *Guide*: Guiding the optimisation solver by locking the satisfactory parts of a solution and leaving the remaining parts for the optimisation solver to change freely.

### 3 DESIGN RECOMMENDATIONS FOR INTERACTIVE OPTIMISATION

Based on this analysis of the problem-solving loop we therefore recommend that an interactive optimisation tool should support the above six tasks. This accords with Endsley’s guidelines for autonomous systems [22] as the task breakdown ensures that the user is in control of the solver’s actions and remains responsible for choosing the final solution. It relies on a client-driven integration model in which the user via the visual interface controls the solver operation [40]. We now give more detailed design recommendations informed by our literature review.

**Recommendation 1: Provide appropriate visual representations of solutions & constraints.** Representations of solutions and constraints should, where possible, make use of visualisations currently used by domain experts so that they can understand and make sense of them. Importantly, the visual representations of solutions and constraints should be tightly coupled. Specifically, when presenting a solution the constraints also need to be represented in order to help users make sense of the solution and to better understand the solution space immediately around it. It may look obvious to represent both solutions and constraints, but on occasion this has been overlooked by optimisation researchers [20].

This recommendation accords with showing contextually relevant information [1], providing effective data representations [21] and use of understandable visualisations [43]. This recommendation is central to supporting interactive optimisation. Without appropriate visual representations of solutions, *Evaluating/Comparing* solutions and *Guiding* the solver will be very difficult. Without appropriate visual representations of constraints, *Improving* a solution and *Refining* an optimisation model is difficult to achieve.

**Recommendation 2: Support user modification of the optimisation model.** The interactive tool should allow users to modify existing constraints, parameters and objective function. Furthermore they should be able to specify new constraints and add them to the optimisation model.

This recommendation supports the *Setup Parameters* and *Refine* steps in the model-defining loop. It provides a semantically rich way of interacting with the solver [21]. Importantly, being able to refine the optimisation model can produce more realistic and satisfying solutions by reducing the gap between the optimisation model and the real-world problem. It also supports dynamic optimisation if the optimisation problem changes.

However, the ability to add constraints is limited by the solver’s capabilities. For instance, the user cannot add non-linear constraints to a linear solver. Furthermore, it can be difficult for non-experts in optimisation to understand abstract constraints. Thus the kinds of constraints that can be added need to be understandable by a domain expert and appropriately represented (Recommendation 1). Constraint acquisition might be a useful approach to elicit constraints from the user [3, 7, 19].

**Recommendation 3: Allow direct manipulation of solutions.** Users should be able to directly manipulate a solution and make whatever changes they like to the values assigned to the decision variables and immediately see the impact on the objective function.

This recommendation allows users to *Improve* an existing solution. Direct manipulation and immediate feedback encourages user engagement [21, 22] and supports efficient correction [1]. Being able to manipulate and change a solution allows users to *Evaluate* its quality in their own way. With the support of appropriate representations of solutions and constraints, the user can explore the trade-offs between solutions, which is helpful for solution evaluation as well as perhaps finding a more satisfying solution. Prior research has found that direct manipulation allows the user to build an appropriate level of trust in the constraint solver [36].

**Recommendation 4: Provide a gallery of solutions.** Multiple solutions should be collected in a gallery. This should provide an interface to explore the solutions and allow users to manually and automatically (re)order solutions based on, say, the value of the objective function.

The solution gallery supports *Comparison* of multiple solutions, allowing users to create a list of ranked solutions based on their preferences. The solution gallery is designed to encourage the user to develop and evaluate alternative solutions. This is particularly important in multi-criteria applications such as design where the objectives are conflicting and the trade-offs are unclear and impossible to encode in a single objective function [18, 57].

**Recommendation 5: Allow user controlled re-optimisation.** We believe it is necessary to provide users with the ability to *re-solve* part of a solution. Often, users are satisfied with only some parts of a solution. In that case, the system can support users by allowing them to “lock” the values of variables in the good part of the solution, and improve the remaining parts of the solution using the re-optimisation. Re-solving provides a semantically rich way of interacting with the solver [21] by allowing the users to *Guide* the optimisation solver to explore the search space that they are interested in. Optimisation solvers should make minimal changes to the parts of a solution that are locked and apply any modifications to the unlocked parts without breaking any constraints. This has the benefit of narrowing the search space to the “unlocked” parts of the solution, so that optimisation solvers can find a better solution more quickly. It also provides stability in the solving process, which helps to preserve the user’s mental model [47, 48] and accords with the guideline for AI systems to update and adapt cautiously [1].

**Recommendation 6: Support comparison of solutions.** To support better *Comparison* of solutions the interface should allow two (or more) solutions to be viewed side-by-side. This allows the user to identify the similarities and differences between the solutions. Visual cues may be used to highlight these. Side-by-side comparisons allows users to *Evaluate/Compare* solutions in a more meaningful way. Solutions with equally-good values of the objective function may differ significantly. Seeing and comparing the trade-offs between these solutions is required in order to find the best solution.

**Recommendation 7: Generate diverse solutions.** We believe an ability to automatically generate a new solution that is significantly different to any of the solutions in the gallery, but which is still of high quality with respect to the objective function, will support the user finding better solutions [18, 33]. It supports *(Re)Solving* the optimisation problem to generate a pool of diverse candidate solutions and can help users better understand the problem by seeing possible trade-offs between different criteria. It helps the user to avoid “decision biasing” [22].

The last two recommendations are not directly motivated by the problem-solving loop but rather are general interface recommendations.

**Recommendation 8: Provide feedback on the solving process.** Often, optimisation software gives little feedback about the solving process, or displays feedback at a very detailed, technical level. This can be disquieting as finding a solution can take minutes, even hours or days. If possible provide the user with understandable feedback on what the solver is doing and how long they will have to wait. Doing so will increase trust in the optimisation software [36]. Even something as simple as a progress bar can show users that the optimisation software is still properly functioning.

Feedback on the solving process might include providing an explanation of the result. *Explainability* of algorithms means to justify the

Recommendations	[32]	[45]	[56]	[10]	[30]	[17]	[4]	[52]	[27]	[8]	[12]	[25]	[13]	[50]	[15]
1a. Provide appropriate visual representations of solutions.															
1b. Provide appropriate visual representations of constraints.															
2a. Support user modification of the optimisation constraints.															
2b. Support user modification of the optimisation objectives.															
3. Allow direct manipulation of solutions.															
4. Provide a gallery of solutions.															
5. Allow user controlled re-optimisation.															
6. Support comparison of solutions.															
7. Generate diverse solutions.															
8. Provide feedback on the solving process.															
9. Record solution provenance.															

Fig. 2. Reflections of the proposed design recommendations in representative systems from a variety of application domains. Yellow cells represent full reflection. Partially shaded cells represent partial reflection/exception. White cells indicate no reflection.

decisions made by an algorithm in a form that is understandable by humans, or other algorithms. Many researchers are exploring explainable ML, e.g. [34] as a way increasing transparency and user trust. Explainable optimisation is also a well established area [29, 44]. However, in practice at present only some algorithms support it and the explanations can be difficult for non-expert users to understand.

**Recommendation 9: Record solution provenance.** As discussed in Recommendation 3 and 7, it is important to allow users to directly manipulate solutions and generate new solutions. During this process it is likely that the users will wish to return to previous solutions. This can be facilitated by storing the history of solution manipulations, and presenting the *provenance* of each solution in terms of the steps that led to it. This encourages user interaction [22] as there is no need to be afraid of making a mistake—they can always go back to a previous solution and start over again. In general, support for provenance is important in visual analytic tools [49]. Another component of provenance is the ability to annotate solutions. This can act as a memory aid for the person using the system and also as a way of communicating different aspects of the solution to stakeholders [35].

### 3.1 Reflections of Design Recommendations in Existing Systems

In order to examine the usefulness of the nine recommendations proposed above, we will now examine (retrospectively) whether the recommendations are met by 15 representative systems from the literature [4, 8, 10, 12, 13, 15, 17, 25, 27, 30, 32, 45, 50, 52, 56]. The examination is not exhaustive as there many interactive optimisation systems. We chose 15 recent systems (selected from 2007 onward) to keep the reflection update-to-date and from different application domains to make the reflection as general as possible. The findings are summarised in Fig. 2.

*Recommendation 1* is very well supported and is the most commonly supported recommendation. All of the examined systems have some kind of visual representation of solutions. This is expected because a suitable visual representation of a solution is fundamental to making the results of the system understandable by the end-user.

Most of these visual representations are domain-specific. For example, in the survey paper of the Human-Guided Search project, Klau et al. [32] listed several interactive optimisation applications. A node-link diagram was used in a graph layout problem. A simplified map containing geographic locations of customers was used in a variation of Travelling Salesman Problem (TSP). Similarly, Belin et al. [8] represented a city as a grid, in which different cells represented different elements such as roads and rivers. However, there were two exceptions. Unlike Belin’s approach, Cajot et al. [12] utilised parallel coordinates in an urban planning problem to facilitate stakeholders’ decision making. In machine learning, Kapoor et al. [30] used an adjacency matrix to represent the classification result. Both of these representations are unlikely to be understood by the average end-user.

The majority of the examined systems also have representations of constraints. For instance, Cummings et al. [17] represented impassable terrain such as big rocks as an area coloured in black in a path planning problem. Bailly et al. [4] drew a lock icon to enforce item placements without changes in a problem of application menu design. However, not

all systems visually represented constraints. In the system developed by Brochu et al. [10], user-defined constraints about model parameters for animation design were shown as a form using raw values.

Recommendations 2, 3, 5 and 7 are supported by roughly half of the examined systems. Specifically, for *Recommendation 2*, some of these systems support constraint modification. For example, Klau et al. [32] allowed users to modify the priority constraint associated with each customer in a VRP, which consequently affected re-optimisation. Some systems utilise users’ preference feedback to influence optimisation models and to guide optimisation solvers in producing more satisfactory solutions. This appeared very common in multi-objective optimisation systems and other systems using genetic algorithms (GA).

For *Recommendation 3*, more than half of the examined systems allowed users to make changes to solutions. For instance, Patten et al. [45] allowed users to directly move physical objects on a table-top interface to re-locate telephone towers when solving a telephone tower layout problem.

For *Recommendation 5*, many examined systems provided re-optimisation. Belin et al. [8] proposed an interactive urban city design tool which the optimisation solver simultaneously improves a city design while users are modifying the design. Most other tools however do not support such simultaneous collaboration and require the users and solvers to take turns to work on the problem [45].

For *Recommendation 7*, half of the examined systems offered different solutions for the users to choose from. Thieke et al. [56] evaluated a medical treatment system which allows clinical professionals to explore multiple equally good treatment plans. It becomes more important to generate and present different solutions to users to facilitate the trade-off comparisons between solutions, especially when there exist multiple conflicting objectives in an optimisation problem. This led to better support of this recommendation in multi-criteria objective problems.

Recommendations 4, 6, 8 and 9 were poorly supported and only reflected in only a few of the examined systems. This is perhaps because these recommendations support higher-level analysis and comparison of solutions rather than the more obvious interactions required to construct a solution. Specifically, for *Recommendation 4*, only Cummings et al. [17] included a very basic solution gallery to facilitate comparison of different path plans. For *Recommendation 6*, Brochu et al. [10] provided four viewports with one solution each in a machine-learning-based optimisation system for animation design. For *Recommendation 8*, Coffey et al. [15] simultaneously updated the model design when users were interacting and modifying the current design. Lastly, for *Recommendation 9*, Klau et al. [32] mentioned that users were allowed to backtrack to previous solutions in the Human-Guided Search project, which provides a partial solution history.<sup>1</sup>

This suggests that while some of the recommendations are typically supported by interactive optimisation systems many are not. The obvious question is therefore whether the proposed recommendations actually lead to more usable interactive optimisation tools.

## 4 INTERACTIVE OPTIMISATION SYSTEM FOR VRPTW

In a second evaluation of the recommendations we used them to guide the development of an exemplar interactive optimisation tool and then conducted a user-study to evaluate the tool and the usefulness of the embodied recommendations.

As our exemplar we chose to develop a tool for solving the Vehicle Routing Problem with Time Windows (VRPTW), a variation of the Vehicle Routing Problem (VRP). This is a practically important problem that is challenging to solve, but understandable by the general public. The VRP problem is to schedule several *trucks* to pick up goods from the *home depot* and deliver them to the *customers* who have ordered the goods. This problem aims to minimise the total *cost* of all *truck routes*. We take total distance travelled by all trucks (the sum of all *truck route* lengths) as the *cost*. VRPTW extends VRP by requiring that a truck has to visit a customer to deliver a service within a certain *time window*. To make the problem even more challenging, we allowed

<sup>1</sup>It is possible that other systems also provided some kind of undo but that this was not mentioned in the system description.





Fig. 3. An overview of the interactive optimisation tool showing its three main components: Solution window, Solution history and Solution gallery (highlighted in blue). Sub-components are highlighted in orange. Interactions are highlighted in red (the rest of interactions in Fig. 4).

a different *service period* for each customer, which is the amount of time the vehicle requires at the customer site to deliver the service (e.g., unloading the goods).

#### 4.1 System Design

We implemented our test system as a web application, with the back-end optimisation based on the MiniZinc [41] system and the Gecode [51] constraint programming solver.

We spent considerable time refining and adjusting the web front-end to conform to the proposed recommendations. The tool contains three main components (see Fig. 3), as follows.

**A) Solution window:** In this window solutions may be explored in detail and interactively modified. We designed the solution window to present either a single solution or two solutions shown side-by-side as recommended by *Recommendation 6: Support comparison of solutions*. This supports participants in comparing different solutions, but it also allows participants to compare a single solution before and after re-optimisation to see what has been changed. Each solution in the solution window is shown in its own *solution viewport*, see Fig. 3. Each solution viewport comprises two panels.<sup>2</sup>

The **street map panel** (Fig. 4-right) displays the geographic information about the trucks' *home depot*, *customers* (labelled A–Z) and *truck routes*—each truck's route represented in a different colour.

We first considered a node-link diagram to display a solution with nodes representing customers (at their geographic locations) and a straight-line link representing the truck route travelled between customers. While these diagrams were easy to draw, and reflect the underlying mathematical model used by the constraint solver, we felt that they violated *Recommendation 1: Provide appropriate visual representation of solutions*. In reality, while a node-link diagram may represent the accurate distance or time between nodes, the route a truck travels from one location to another can be complex, with many turns at intersections and roundabouts.

<sup>2</sup>The representation of solutions is similar to that used in [36]. However, the system used in [36] was a minimal implementation designed only to evaluate the effect of solver feedback and solution manipulation on trust and so provided limited support for interaction and no support for most recommendations.

We found a better representation using OpenStreetMap (OSM)<sup>3</sup> for map display and Open Source Routing Machine (OSRM)<sup>4</sup> for realistic routing in road networks, together with Leaflet<sup>5</sup>, a JavaScript map annotation library, to create the street map representation for our tool. Now, the street map view provides an unambiguous visual context of a VRPTW and the routes are represented both realistically and accurately, being faithful to the underlying mathematical model.

The **schedule panel** (Fig. 4-left) shows the time window of each customer and the order of service delivery of each truck. Marey diagrams [58] are used in transport schedules to represent locations of vehicles and the transport routes, which fits well in the VRPTW context. In Marey diagrams time is showed vertically and locations are displayed horizontally. However, to better suit the aspect ratio of the panel we chose to swap these. This also has the advantage of making customer names more visually salient than the time of delivery, reflecting their relative importance.

Thus a schedule has each customer represented by a horizontal line and the horizontal extent of the lines indicates the total time required for the schedule. Customer lines are grouped by the truck which services them. A sequence of line segments in the colour of the truck as used on the map, passes through the customer lines, indicating the truck's progress in making deliveries to customers. Time windows for each customer are represented by a grey rectangle with the position of the left side of this rectangle indicating the start of the window, and the right side indicating the end of the window. The other essential constraint is the service period, i.e. the time required by a truck to deliver a service. We used another rectangle in the colour of the truck to represent a service period. The rectangular representations of the time window and the service period were chosen for two reasons: First, axis alignment facilitates quantitative comparisons of a service period and a time window of a single customer and comparisons of time windows of multiple customers. Second, it visually distinguishes valid from invalid situations. In a valid situation, a service period stays within the range of a time window resulting in the service period

<sup>3</sup><https://www.openstreetmap.org>

<sup>4</sup><https://project-osrm.org>

<sup>5</sup><https://www.leafletjs.com>

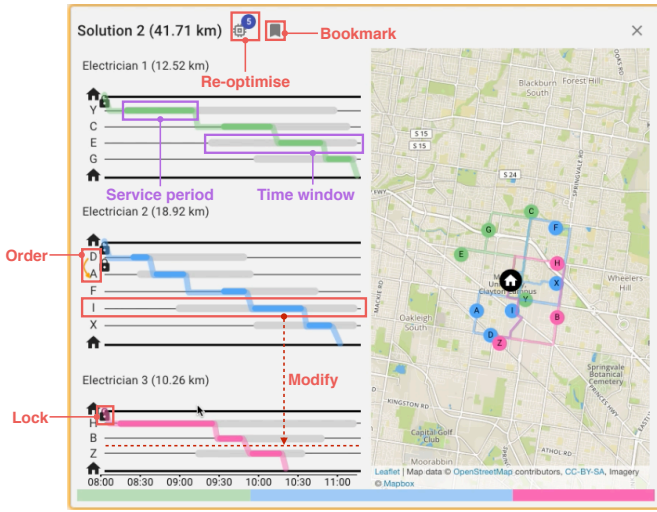


Fig. 4. A detailed view of the *solution viewport* in the solution window including the street map and the schedule. Interactions are highlighted in red. Constraints are marked in purple.

rectangle completely overlapped on top of the time window rectangle. Whereas in an invalid situation, a service period goes beyond the close of a time window resulting in a possible partial overlapping between the service period rectangle and the time window rectangle. Thus this representation conforms to the second part of *Recommendation 1: Provide appropriate visual representation of constraints*.

**B) Solution history:** keeps tracks of the participant's interaction history. We designed a histogram to record all solutions and their provenances according to the *Recommendation 9: Record solution provenance*.

**Objective histogram** represents each solution as a rectangular bar whose height indicates the value of the solution's objective function (total distance travelled). Whenever constraints of a solution change or the solution is re-optimised, we draw a new bar for the new solution. We draw an arc between two solutions if the solution on the right is derived from that on the left. Selecting a history bar brings the details of the corresponding solution into the solution window.

**C) Solution gallery:** supports participants managing solutions so as to finally select the best solution out of many, following *Recommendation 4: Provide a gallery of solutions*. The gallery supports annotations allowing participants to edit solution names (*Recommendation 9: Record solution provenance*).

**Mini-solution** provides a simplified view of a solution by showing vehicle routes without the street map and including the objective value and the solution name.

The tool initialises the solution history and gallery with several different solutions to give users flexibility to choose between different starting solutions. The solutions are chosen to be diverse. We measured diversity between solutions using a diversity matrix, which calculated the routing sequence differences between one solution and all the others. We select solutions with maximum diversity between the new computed solution and any previously computed solutions. Providing such a selection of starting points meets *Recommendation 7: Generate diverse solutions*.

## 4.2 System Interaction

Interactions fall into two groups: **optimisation-model interaction** and **problem-solving interaction**. There are two optimisation-model interactions:

**Lock** allows the user to fix a specific customer to a specific truck. A lock is drawn beside the customer to represent this constraint. This allows us to specify, for instance, that a customer requires the use of a particular truck. Locking a customer to a truck does not constrain the time of service by the truck.

**Order** fixes the relative order in which two customers of the same truck are to be serviced. A directed arc line is drawn from one customer

to another to represent this constraint. It is useful in pick-up and delivery problems, where the goods to be delivered to a customer may be located at another customer's location. Thus, the goods from the second customer need to be picked up first and then delivered to the first customer.

These constraints, supported by appropriate visual representations (*Recommendation 1: Provide appropriate visual representation of constraints*), give users the flexibility to modify the optimisation models, *Recommendation 2: Support user modification of the optimisation model*.

There are five problem-solving interactions. These allow the user to create new solutions. **Modify** allows users to re-assign customers within the same truck or between two trucks via drag-and-drop. This conforms to *Recommendation 3: Allow direct manipulation of solutions*.

**Re-optimize** allows the user to collaboratively solve an optimisation problem with the optimisation solver with the control of when and what to re-optimize. This is a reflection of the *Recommendation 5: Allow user controlled re-optimisation*. In particular, users can freeze a part or multiple parts of a solution using the optimisation-model interactions described above, and then re-optimize the solution by clicking the re-optimize button located at the very top of the solution viewport within the solution window.

Before the optimisation solver spits out a new solution, the re-optimize button is replaced by a spinning wheel to indicate the ongoing re-optimisation process. This is in accord with *Recommendation 8: Provide feedback on the solving process*.

**Load** allows users to choose a solution from either the solution history or the solution gallery and to display the solution on either the left or the right solution viewport in the solution window by clicking the load button when hovering over the target solution.

**Bookmark** enables users to mark and store a solution in the solution gallery by clicking the bookmark button when hovering over a solution in either the solution window (either viewport) or the solution history.

**Rename** allows users to change the name of a solution. The solution name can become a useful annotation allowing for organising and finding candidate solutions.

## 5 USER STUDY

Our user study of the interactive optimisation tool was designed to evaluate: (1) Usability of the tool; (2) The tool's ability to support unforeseen scenarios (flexibility); (3) The tool's support for the recommendations and usefulness of these recommendations as evidenced by its use.

### 5.1 Study Design

**Participants:** We recruited 10 participants: 7 females and 3 males. They were PhD students from our university and employees from other organisations with a bachelor's degree or higher. As end-users of interactive optimisation tool are more likely to be non-experts rather than optimisation experts, we only selected participants who did *not* have computer science or optimisation expertise. All participants had a normal or corrected-to-normal vision and were without any colour vision impairment.

**Procedure:** Participants used our interactive optimisation tool throughout the study. The application domain presented to participants was that of a company that sends electricians (in trucks) to customers. Each customer requires a certain fixed amount of time for the service, and this has a time window. All trucks start from a central depot and return to the depot after servicing their customers.

The study took roughly one hour and ten minutes on average. The study was run on a MacBook Pro notebook with a 2.6 GHz Intel i5 processor and a 13-inch screen. It had three parts:

(a) *Training:* This provided an introduction to the application domain, the solution and constraint representation, interface components (solution window, solution history and solution gallery) as well as user interactions (*Modify* a solution by re-assigning customers; *Add & remove* a customer *lock* constraint; *Add & remove* a customer *order* constraint; *Re-optimize* a solution when constraints have been modified;

R1. Provide appropriate visual representations of solutions & constraints.
Q1-1 I found the visual representation of a solution to the routing problem to be intuitive (the route on the map and the electrician schedule).
Q1-2 I found the visual representation of the constraints to be intuitive (the locks and the arrows).
Q1-3* Could these be improved? How?
R2. Support user modification of the optimisation model.
Q2-1 I found the system provided enough flexibility to model customer and operation requirements by adding new constraints (locks and arrows).
Q2-2* Are there other constraints you would like the system to support?
R3. Allow direct manipulation of solutions.
Q3-1 When moving customers around within a single electrician's schedule or between electricians I found it useful to see immediately how this affects the solution (distance travelled, feasibility).
R4. Provide a gallery of solutions.
Q4-1 I found the solution gallery useful.
Q4-2* Could this be improved? How?
R5. Allow user controlled re-optimisation.
Q5-1 I found it useful to re-optimize after changing constraints.
R6. Support comparison of solutions.
Q6-1 I found it useful to see two solutions side-by-side.
Q6-2* Could this be improved? How?
R7. Generate diverse solutions.
Q7-1 I found it useful to start with a number of different solutions in the solution gallery.
R8. Provide feedback on the solving process.
Q8-1 I found the spinning wheel provided sufficient feedback while the system is producing a new solution.
Q8-2* Would you like more feedback?
R9. Record solution provenance.
Q9-1 I found it useful to go back to previous solutions.
Q9-2 I could understand the history of the solutions.
Q9-3* Could these be improved? How?

Fig. 5. (Q)uestions about interface features and their usefulness based on each (R)ecommendation. Questions marked with \* are open-ended questions.

Load a solution to the solution window; *Bookmark* a solution to the solution gallery; *Rename* a solution in the solution gallery.)

(b) *Problem-solving*: Participants were then asked to solve an optimisation problem in four different scenarios. The first scenario was the tool's intended use-case while the other three were designed to evaluate its flexibility in situations it wasn't explicitly designed to handle. At the end of each scenario, participants were asked to explain their strategies for solving the problem.

In *Scenario 1* participants were asked to find a solution minimising the total travelling distance with no constraints apart from satisfying customer time-windows. This was the tool's intended use-case. In Part 1 of this scenario, participants could manually improve the solutions but re-optimisation was not allowed, then in Part 2 they were allowed to use re-optimisation. This was to ensure that they experimented with both manual manipulation and re-optimisation. Participants were given a maximum of 5 minutes for each part of the scenario.

*Scenario 2* asked participants to solve a question of similar difficulty to Scenario 1. However, we introduced another objective, asking participant's to minimise travel distance but also balance the workload of each truck so that they service close to the same number of customers. There was a 5-minute time limit for this Scenario.

*Scenario 3* asked participants to find a solution satisfying a list of customer requests. The two objectives remained the same as in Scenario 2. Most of the requests could be satisfied by adding either a customer lock constraint or a customer order constraint. For instance, one request was "Customer X requires the service from Electrician 2". However, there was a more difficult request specifying that one customer should not be serviced by a particular truck. This was difficult because none of the two customer-related constraints can directly encode it. Because of the difficulty, participants had 10 minutes for this scenario.

In *Scenario 4*, participants were asked to revise the schedule as the result of an emergency in which one of the trucks has broken down. The two objectives remained the same as in Scenario 2. This was challenging because not only did participants need to re-assign the remaining customers of the broken truck but also participants had to ensure that all services delivered before the truck break-down time were unchanged, as these events were now in the past. There were no constraints that could directly achieve this, which made this problem very difficult. For this reason, we gave participants 15 minutes.

(c) *System evaluation*: After completing the scenarios participants were asked to evaluate the system *interface* and *usability*. Participants were first asked to rank different aspects of the interface using a Likert scale from 1 to 5. The questions were designed to evaluate the usefulness of these features and hence the usefulness of the underlying recommendations. We decided not to ask participants to evaluate the recommendations directly because the recommendations were rather abstract and difficult for non-experts to understand. The questions can be found in Fig. 5. We used the standard System Usability Scale (SUS) [11], a 10-item questionnaire to measure the usability of our interface. Again participants were asked to rank each item using the same 5-point Likert scale. We did not modify the questions in the SUS questionnaire. However, we did provide extra context for the first and the last questions in order to clarify what was being asked. For example, the first question was 'I think that I would like to use this system frequently.' We provided the extra context 'Assume that you need to solve similar vehicle routing problems as a part of your job.'

We prepared four problem instances of the VRPTW with similar complexity for the four scenarios (3 trucks, 12 customers). For the first three scenarios, we pre-computed 3 diverse solutions that were stored in both the solution history and the solution gallery. The last scenario, started with a single solution. In terms of the solution quality, the first two problem instances contained *poor* solutions which were 30% worse than the best solution. The last two problems contained *best* or close-to-best (with 5% solution quality margin) solutions. We did this because we thought this would encourage participants to explore different solutions in the first two scenarios as finding an improved solution was relatively easy.

## 5.2 Results & Discussion

Recall we wished to evaluate: (1) Usability of the tool; (2) Its flexibility to support unforeseen scenarios; (3) Support for the recommendations and their usefulness.

**Usability**: At the end of the study participants were asked to evaluate the usability of the tool using the SUS (see Fig. 6). We used the standard approach to calculate the overall system usability score: convert the ratings to values from 0 to 4 (4 being the most positive response), subtract the converted rating of each even-numbered questions from 4 because these questions were in a negative tone, sum these converted ratings of all odd- and even-numbered questions, then multiply the summed rating by 2.5 to give a score out of 100.

The highest score was 97.5, and the lowest score was 77.5. Except for the lowest score, all other scores were equal to or greater than 87.5. On average, the SUS score of our tool was 91, which is far better than the acceptance threshold value of 70 [5]. Generally, a score above 80 means that the tool is of excellent usability, and is acceptable to end-users. Thus we can be reasonably confident that usability issues did not impact the use of different features.

**Flexibility**: The different scenarios were designed to explore the tool's flexibility. All participants finished both parts of Scenario 1. 8 out of 10 participants developed a better solution in Part 1. However, none of them found an optimal solution, i.e. with the minimum total travel distance. In general, participants tried to group close-by customers and service them with the same truck to reduce the total travelling distance.




Fig. 6. Participants' ratings of SUS questionnaire.



In Part 2, an optimal solution was discovered by all participants using re-optimisation.

Most participants found very good solutions to Scenario 2. These solutions were better than the starting solutions in two aspects: not only was the total distances shorter than any of the pre-computed diverse solutions, but also participants' developed solutions that were perfectly balanced with 4 customers per truck. Having the second objective to also balance trucks' workloads affected participants' problem-solving strategies.

Specifically, 7 participants explained that they balanced the workload first by re-assigning customers, and then minimised the total travelling distance using re-optimisation. In contrast, the other 3 participants tried to minimise the distance first using re-optimisation, and then to manually re-assign customers for workload balancing. One participant had an interesting strategy. (S)he primarily focused on re-assigning customers with long time windows because these customers were easier to fit in the schedule compared to short time-window customers. Thus, such customer re-assignments were more tolerant of time window violations and it was easier to fix the solution when violations occurred.

Scenario 3 introduced customer requests. All participants managed to find high quality solutions satisfying the requests and balancing the two objectives. All participants added constraints to enforce the majority of the customer requests and then tried to balance the two objectives using the same approaches as in Scenario 2. Most used lock to ensure that the customer who did not like a particular electrician remained assigned to one of the other electricians.

Scenario 4 asked participants modify the schedule after a truck breakdown. 9 out of 10 participants developed a valid solution. All participants agreed that the emergency should be handled first, i.e. first find a solution that no longer used the broken truck but that preserved the schedule for the time before the breakdown. To do so they re-assigned customers from the broken truck to the other two trucks. When doing so they also attempted to balance the workloads. They then attempted to use re-optimisation to reduce the overall travel distance. However, this typically produced an invalid solution as the optimisation solver could assign a customer back to the broken truck if the customer was unlocked.

After realising this, some participants decided to lock all 12 customers to prevent unwanted customer re-assignments before re-optimising. This only worked for one fortunate participant. The other participants found that the optimisation solver could not re-optimize the current solution due to conflicting constraints.

As a result, most participants developed their valid solutions through an iterative process, in which they re-assigned customers first before utilising re-optimisation to improve the travelling distance. Once re-optimisation was done, they re-assigning customers before the next round of re-optimisation. This process terminated when a valid solution was found. However, three participants solved the problem using a purely manual approach to re-assign customers without relying on re-optimisation to find a solution.

Thus we see that, as hoped, the tool allowed the participants to flexibly handle a variety of scenarios which the tool was not explicitly designed to support. This relied upon the ability to add/modify/remove customer constraints, re-optimize taking into account modified constraints, and to manually modify a solution.

**Support for the recommendations and their usefulness:** Fig. 8 shows the number of participants who used various features of the interface in each scenario and Fig. 7 summarises participants response to the questions about the system interface. These, together with the participant strategies detailed above, allow us to gauge how well the tool supported each recommendation and at least to some extent the usefulness/importance of the underlying recommendations. We first detail the recommendations with strong support.

**Recommendation 1:** This was well supported by the tool. Most participants gave 4 out of 5 for the visual representations of the solution (Q1-1) and constraints (Q1-2) (see Fig. 5). Participants agreed that both the schedule and street map were clean and well designed. One participant said "The schedule looks very straightforward, and from the map, I can clearly see customers as well as their [assigned trucks]."

Another commented "It is very easy to understand what is happening," implicitly suggesting that the visualisation is useful. Participants did suggest two improvements (Q1-3): to add route arrows on the street map to indicate travel direction, and to place the customer lock constraint label in front of the customer letter. While participant comments did not directly address the importance of this recommendation, understandable solution and constraint representation underlie Recommendations 1-4 and 6, most of which we found to be important.

**Recommendation 2:** This was also well supported by the tool. On average, participants ranked 4.5 for flexibility (Q2-1). Our observations of participants suggested they understand both customer lock and order constraints, and they could effectively use these to model more complex requirements. From Fig. 8 we see that they were used by at least half of the participants in each scenario except for Scenario 1 Part 1, and by all in Scenario 3. This argues strongly for the importance of this recommendation to provide flexibility in unforeseen situations.

When participants were asked if there were other constraints they would like the system to support (Q2-2), some suggested that it would be nice to have a truck lock constraint which locks all customers serviced by the same truck and also prevents other customers being assigned to this particular truck during re-optimisation. Some also suggested that apart from the existing "relative" customer lock which only guarantees a customer to be serviced by a specific truck, it would be nice to have an "absolute" customer lock so that once a customer is locked, the service order of this particular customer is also locked.

**Recommendation 3:** This was very well supported. The majority of the participants ranked 5 for solution manipulations and immediate feedback after manipulations (Q3-1). All participants agreed that it was necessary to provide such immediate feedback when changes are made. One participant commented "I think it (manipulating a solution and providing feedback) is very useful. Because not only I can see how distance (one of the objectives) has changed, but also I can see whether the solution has any violations highlighted in red." Manual adjustment of solutions was performed by all participants in all scenarios, strongly suggesting the importance of this recommendation.

**Recommendation 5:** This was very well supported with all but one participant finding re-optimisation very useful (Q5-1) and the other finding it useful. One participant commented "This is extremely useful because I feel like I do not need to modify the solution again after re-optimisation." Re-optimisation was used by all participants in three scenarios and by most in Scenario 4.

Some participants suggested that it would be nice to clearly see the changes between solutions. Two participants advised demonstrating the solution changes before and after re-optimisation with one saying "After re-optimisation, the orders of customers may have changed drastically. Therefore it would be great to show what and where the changes are." One participant commented "If it can show animations to indicate how does one solution transform into another, that would be useful because it could reduce my memory load and remind me what has changed." This should be explored in future work.

**Recommendation 8:** There was qualified support for this. Most participants believed having the spinning wheel during re-optimisation provided sufficient feedback about the ongoing solving process (Q8-




Fig. 7. Participants' ratings of interface questions.

		Lock	Order	Modify	Re-optimize	Load	Bookmark	Rename
Scenario 1	Part 1	3	0	10	4	0	0	0
	Part 2	5	1	5	10	4	1	0
Scenario 2		5	0	10	10	2	0	0
Scenario 3		10	10	10	10	8	0	0
Scenario 4		7	3	10	7	6	0	0
		9-10	7-8	5-6	3-4	1-2	0	0

Fig. 8. Participants use of different kinds of interaction in each scenario.



1). In particular, one participant commented *"It is a must-have. It is rather a psychological effect telling me that the re-optimisation is still running."* Another three participants made similar comments. However, three participants thought the spinning wheel was not useful. This is because they were too used to it. One participant said *"It is so common that every system has this (a spinning wheel or something similar to indicate task progress)."*

When participants were asked whether more feedback is needed (Q8-2), several suggestions were provided. Specifically, one participant mentioned a progress bar and commented *"I think it (the spinning wheel) is enough since it (the re-optimisation) is fast. If it (the re-optimisation) is slow, you might want to consider something like a progress bar."* Another participant suggested having a pop-up message to briefly summarise all objective values.

Overall, this suggests that while providing feedback on the solver's progress is generally appreciated, more research is needed into how to do so. One must be careful not to provide feedback that leads to over trust [36].

**Recommendation 9** was very well supported. Every participant believed that it was essential to be able to go back to previous solutions (Q9-1). For instance, one participant said *"It is really useful because you may make a mistake, and you may also not remember the previous solution."* All participants also agreed that the history of the solutions is easy to understand (Q9-2). Participant's behaviour also showed that it was relatively common to load a solution from the history in many scenarios (see Fig. 8).

**Recommendations involving multiple solutions received less support.** Recommendations 4, 6 and 7 were generally not as well supported by participant feedback, as follows.

**Recommendation 4:** In particular, we were surprised to find that participants did not find the solution gallery to be useful with an average ranking of 3 (Q4-1). Only one participant used the *bookmark* interaction, and none used the *rename* interaction though five participants used the *load* interaction. When participants were asked how to improve the solution gallery (Q4-2), one suggested *"If there is a way to [automatically] re-order solutions based on quality, it would be great, for instance, to place the best solution at the very top and the worst at the very bottom."*

**Recommendation 6:** While this was supported with all participants agreeing that seeing two solutions side-by-side was useful (Q6-1). In practice we did not observe participants spending a lot of time comparing solutions. Instead they appeared to focus on one solution and try and improve it. Considering participants' suggestions on the side-by-side layout (Q6-2), more than half of the participants thought having two solutions presented simultaneously would be sufficient. For instance, one participant commented *"I think two is enough. With more than two solutions, I won't be able to simultaneously look at them all."* However, one participant said *"If you can provide options [of how many solutions to be presented], depending on the size of monitors, they can make it as many [solutions] as they want to. But I think two is enough."*

**Recommendation 7:** This had more support with most participants agreeing that it is beneficial to have several different solutions in the solution gallery to choose from (Q7-1). In Scenario 3, eight participants conducted a quick examination of all solutions in the gallery to try to find the solution satisfying the greatest number of customer requests. They then used this as their starting point. However, our observations revealed that after an initial examination, most participants focused on one solution and did not switch to another solution. Nonetheless, in Scenarios 4, we did see three participants switch to another solution when they became stuck. The situation improved after switching, resulting in better solution performance.

We suspect that one reason that many participants did not use the gallery or explore multiple solutions was the enforced time limit in each scenario. We observed that some participants utilised the maximum amount of time in several scenarios and one participant stated that *"If you do not put the time limit [for each scenario], it (the solution gallery) would be useful. With the time limit, I have to hurry up and focus on only one solution."* When asked *"What if there was no time limit, then how would you use it?"* they responded *"I would like to store*

*all good solutions in it (the solution gallery) and do comparisons afterwards."* Another participant said *"Without time limitation, I would like to develop two solutions with one solution per viewport and choose the better one in the end."* A further participant made a similar comment.

Lack of experience with the tool may also have contributed. One participant commenting *"Maybe if I use this software more, I would have more experience [using the solution gallery]."*

Another part of the reason is that participants felt the problems were simple enough that they did not feel the need to explore multiple solutions. One said *"Actually I only deal with one solution at a time. But if there are a lot of solutions, then maybe that (the solution gallery) helps."* Another one responded *"It might be more useful when the problem becomes more complex with more solutions provided."*

We conjecture that a solution gallery seeded with diverse solutions and the ability to compare solutions would be more useful in more open-ended tasks such as design involving complex multi-criteria optimisation. This would accord with Dayama et al. [18] who found designers liked to have a gallery of automatically suggested solutions. Such a gallery should support ranking based on multiple criteria. This was supported by the suggestions of two participants who mentioned customer distributions within trucks. Specifically, one said *"It would be better to include the workload balance (customer distribution) information of each electrician [for each solution in the solution gallery]."*

## 6 CONCLUSION

We have presented a comprehensive set of design recommendations (R1-9) for interactive optimisation systems. To the best of our knowledge these are the first such recommendations. An examination of 15 representative systems from the literature revealed that none supported all of the recommendations. Some recommendations—appropriate visual representation and modifiable constraints—were well supported, some were supported by about half of the systems—manual modification of a solution, user controlled re-optimisation and generation of diverse solutions—while the remaining recommendations—solution gallery, side-by-side solution comparison, solver feedback and record of solution provenance—were rarely supported.

We then evaluated the recommendations using the vehicle routing problem with time windows (VRPTW) as an exemplar application. We built an interactive tool for solving these problems that was informed by the recommendations. Ten participants then used this system to solve a variety of routing problems. These showed that the tool was very flexible and could be used to solve variants of VRPTW it was not originally designed for. Participants found the tool highly usable and, based on user feedback and behaviour, we found strong support for most of the underpinning recommendations.

The recommendations supporting exploration of multiple solutions (R4 and R6) received less support in our observations. We conjecture that this is likely to reflect the limited time allowed in the study; the simplicity of the objective function; and the nature of the tasks they were required to perform. Further studies are required to investigate this, ideally with longitudinal deployment of interactive optimisation into a problem-solving workflow.

Of course, in general, more case studies are required in order to confirm the general applicability of the recommendations. Nonetheless we believe that our recommendations are a valuable addition to our understanding of how to design interactive optimisation systems as the features they suggest are useful but many are not commonly found in existing systems.

## ACKNOWLEDGMENTS

We thank Agilent Technologies Inc. for their support through a Thought Leader Award and the support of Data 61, CSIRO which is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre for Excellence Program. We also thank the reviewers for their insightful comments and suggestions.

## REFERENCES

- [1] S. Amershi, D. Weld, M. Vorvoreanu, A. Fourney, B. Nushi, P. Collisson, J. Suh, S. Iqbal, P. N. Bennett, K. Inkpen, et al. Guidelines for human-ai interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*, pp. 1–13, 2019.
- [2] D. Anderson, E. Anderson, N. Lesh, J. Marks, B. Mirtich, D. Ratajczak, and K. Ryall. Human-guided simple search. In *AAAI/IAAI*, pp. 209–216, 2000.
- [3] R. Arcangioli, C. Bessière, and N. Lazaar. Multiple Constraint Acquisition. In *IJCAI: International Joint Conference on Artificial Intelligence*, pp. 698–704. New York City, United States, July 2016.
- [4] G. Bailly, A. Oulasvirta, T. Kötzing, and S. Hoppe. Menuoptimizer: Interactive optimization of menu systems. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*, pp. 331–342. ACM, 2013.
- [5] A. Bangor, P. T. Kortum, and J. T. Miller. An empirical evaluation of the system usability scale. *Intl. Journal of Human-Computer Interaction*, 24(6):574–594, 2008.
- [6] C. Barnhart and G. Laporte, eds. *Handbooks in Operations Research and Management Science: Transportation*, vol. 14. North Holland, 2006.
- [7] N. Beldiceanu and H. Simonis. A constraint seeker: Finding and ranking global constraints from examples. In *International Conference on Principles and Practice of Constraint Programming*, pp. 12–26. Springer, 2011.
- [8] B. Belin, M. Christie, and C. Truchet. Interactive design of sustainable cities with a distributed local search solver. In *Integration of AI and OR Techniques in Constraint Programming*, pp. 104–119. Springer, 2014.
- [9] J. M. Betts, C. Mears, H. M. Reynolds, G. Tack, K. Leo, M. A. Ebert, and A. Haworth. Optimised robust treatment plans for prostate cancer focal brachytherapy. In S. K. et al., ed., *ICCS 2015*, vol. 51 of *Procedia Computer Science*, pp. 914–923. Elsevier, 2015.
- [10] E. Brochu, T. Brochu, and N. de Freitas. A Bayesian interactive optimization approach to procedural animation design. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pp. 103–112. Eurographics Association, 2010.
- [11] J. Brooke et al. Sus-a quick and dirty usability scale. *Usability evaluation in industry*, 189(194):4–7, 1996.
- [12] S. Cajot, N. Schüler, M. Peter, A. Koch, and F. Maréchal. Interactive optimization with parallel coordinates: exploring multidimensional spaces for decision support. *Frontiers in ICT*, 5:32, 2019.
- [13] L. Caldas and L. Santos. Painting with light: An interactive evolutionary system for daylighting design. *Building and Environment*, 109:154–174, 2016.
- [14] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, and C. Tominski. Characterizing guidance in visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):111–120, 2016.
- [15] D. Coffey, C.-L. Lin, A. G. Erdman, and D. F. Keefe. Design by dragging: An interface for creative forward and inverse design with simulation ensembles. *IEEE transactions on visualization and computer graphics*, 19(12):2783–2791, 2013.
- [16] C. Collins, N. Andrienko, T. Schreck, J. Yang, J. Choo, U. Engelke, A. Jena, and T. Dwyer. Guidance in the human-machine analytics process. *Visual Informatics*, 2(3):166–180, 2018.
- [17] M. L. Cummings, J. J. Marquez, and N. Roy. Human-automated path planning optimization and decision support. *International Journal of Human-Computer Studies*, 70(2):116–128, 2012.
- [18] N. R. Dayama, K. Todi, T. Saarelainen, and A. Oulasvirta. Grids: Interactive layout design with integer programming. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2020.
- [19] L. De Raedt, A. Passerini, and S. Teso. Learning constraints from examples. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [20] P. Deransart and M. V. Hermenegildo. *Analysis and visualization tools for constraint programming: constraint debugging*. Number 1870. Springer Science & Business Media, 2000.
- [21] J. J. Dudley and P. O. Kristensson. A review of user interface design for interactive machine learning. *ACM Transactions on Interactive Intelligent Systems (TiIS)*, 8(2):1–37, 2018.
- [22] M. R. Endsley. From here to autonomy: lessons learned from human-automation research. *Human factors*, 59(1):5–27, 2017.
- [23] K. Z. Gajos, D. S. Weld, and J. O. Wobbrock. Automatically generating personalized user interfaces with supple. *Artificial Intelligence*, 174(12-13):910–950, 2010.
- [24] S. Goodwin, C. Mears, T. Dwyer, M. G. de la Banda, G. Tack, and M. Wallace. What do constraint programming users want to see? Exploring the role of visualisation in profiling of models and search. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):281–290, 2017.
- [25] J. Hakonen, K. Miettinen, and K. Sahlstedt. Wastewater treatment: New insight provided by interactive multiobjective optimization. *Decision Support Systems*, 51(2):328–337, 2011.
- [26] F. S. Hillier. *Introduction to operations research*. Tata McGraw-Hill Education, 2012.
- [27] A. Holzinger, M. Plass, M. Kickmeier-Rust, K. Holzinger, G. C. Crişan, C.-M. Pintea, and V. Palade. Interactive machine learning: experimental evidence for the human in the algorithmic loop. *Applied Intelligence*, 49(7):2401–2414, 2019.
- [28] C. V. Jones. Visualization and optimization. *ORSA Journal on Computing*, 6(3):221–257, 1994.
- [29] U. Junker. QUICKXPLAIN: preferred explanations and relaxations for over-constrained problems. In D. L. McGuinness and G. Ferguson, eds., *Proceedings of the Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference on Innovative Applications of Artificial Intelligence, July 25-29, 2004, San Jose, California, USA*, pp. 167–172. AAAI Press / The MIT Press, 2004.
- [30] A. Kapoor, B. Lee, D. Tan, and E. Horvitz. Interactive optimization for steering machine classification. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1343–1352. ACM, 2010.
- [31] D. A. Keim, J. Kohlhammer, G. Ellis, and F. Mansmann. *Mastering the Information Age: Solving Problems with Visual Analytics*. Eurographics Association, 2010.
- [32] G. W. Klau, N. Lesh, J. Marks, and M. Mitzenmacher. Human-guided search. *Journal of Heuristics*, 16(3):289–310, 2010.
- [33] J. Koch, A. Lucero, L. Hegemann, and A. Oulasvirta. May ai? design ideation with cooperative contextual bandits. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2019.
- [34] T. Kulesza, M. Burnett, W.-K. Wong, and S. Stumpf. Principles of explanatory debugging to personalize interactive machine learning. In *Proceedings of the 20th international conference on intelligent user interfaces*, pp. 126–137, 2015.
- [35] J. Liu, T. Dwyer, K. Marriott, J. Millar, and A. Haworth. Understanding the relationship between interactive optimisation and visual analytics in the context of prostate brachytherapy. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):319–329, 2018.
- [36] J. Liu, K. Marriott, T. Dwyer, and G. Tack. Increasing user trust in optimisation through feedback and interaction. *submitted and included in supplementary material*.
- [37] J. Matejka, M. Glueck, E. Bradner, A. Hashemi, T. Grossman, and G. Fitzmaurice. Dream lens: Exploration and visualization of large-scale generative design datasets. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2018.
- [38] D. Meignan, S. Knust, J.-M. Frayret, G. Pesant, and N. Gaud. A review and taxonomy of interactive optimization methods in operations research. *ACM Transactions on Interactive Intelligent Systems (TiIS)*, 5(3):17, 2015.
- [39] K. Miettinen. Survey of methods to visualize alternatives in multiple criteria decision making problems. *OR spectrum*, 36(1):3–37, 2014.
- [40] T. Mühlbacher, H. Piringer, S. Gratzl, M. Sedlmair, and M. Streit. Opening the black box: Strategies for increased user involvement in existing algorithm implementations. *IEEE transactions on visualization and computer graphics*, 20(12):1643–1652, 2014.
- [41] N. Nethercote, P. J. Stuckey, R. Becket, S. Brand, G. J. Duck, and G. Tack. MiniZinc: Towards a standard CP modelling language. In C. Bessière, ed., *Principles and Practice of Constraint Programming - CP 2007, 13th International Conference, CP 2007, Providence, RI, USA, September 23-27, 2007, Proceedings*, vol. 4741 of *Lecture Notes in Computer Science*, pp. 529–543. Springer, 2007.
- [42] J. Nocedal and S. Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [43] G. Nodalo, J. M. Santiago III, J. Valenzuela, and J. A. Deja. On building design guidelines for an interactive machine learning sandbox application. In *Proceedings of the 5th International ACM In-Cooperation HCI and UX Conference*, pp. 70–77, 2019.
- [44] B. O’Callaghan, B. O’Sullivan, and E. C. Freuder. Generating corrective explanations for interactive constraint satisfaction. In P. van Beek, ed.,

*Principles and Practice of Constraint Programming - CP 2005*, pp. 445–459. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.

- [45] J. Patten and H. Ishii. Mechanical constraints as computational constraints in tabletop tangible interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 809–818. ACM, 2007.
- [46] P. Pirolli and S. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, vol. 5, pp. 2–4. McLean, VA, USA, 2005.
- [47] H. C. Purchase, E. Hoggan, and C. Görg. How important is the “mental map”?—an empirical investigation of a dynamic graph layout algorithm. In *International Symposium on Graph Drawing*, pp. 184–195. Springer, 2006.
- [48] H. C. Purchase and A. Samra. Extremes are better: Investigating mental map preservation in dynamic graphs. In *International Conference on Theory and Application of Diagrams*, pp. 60–73. Springer, 2008.
- [49] E. D. Ragan, A. Endert, J. Sanyal, and J. Chen. Characterizing provenance in visualization and data analysis: an organizational framework of provenance types and purposes. *IEEE transactions on visualization and computer graphics*, 22(1):31–40, 2015.
- [50] D. Schneeweiss and P. Hofstedt. Fdconfig: a constraint-based interactive product configurator. In *Applications of Declarative Programming and Knowledge Management*, pp. 239–255. Springer, 2011.
- [51] C. Schulte, M. Lagerkvist, and G. Tack. Gecode. Software and documentation available from <https://www.gecode.org>.
- [52] C. Schumann, C. Rieder, S. Haase, K. Teichert, P. Süß, P. Isfort, P. Bruners, and T. Preusser. Interactive multi-criteria planning for radiofrequency ablation. *International journal of computer assisted radiology and surgery*, 10(6):879–889, 2015.
- [53] J. P. Shim, M. Warkentin, J. F. Courtney, D. J. Power, R. Sharda, and C. Carlsson. Past, present, and future of decision support technology. *Decision support systems*, 33(2):111–126, 2002.
- [54] G. Simonin, C. Artigues, E. Hebrard, and P. Lopez. Scheduling scientific experiments for comet exploration. *Constraints*, 20(1):77–99, 2015.
- [55] M. Stølevik, T. E. Nordlander, A. Riise, and H. Frøyseth. A hybrid approach for solving real-world nurse rostering problems. In J. H. Lee, ed., *CP 2011*, vol. 6876 of *Lecture Notes in Computer Science*, pp. 85–99. Springer, 2011.
- [56] C. Thieke, K.-H. Küfer, M. Monz, A. Scherrer, F. Alonso, U. Oelfke, P. E. Huber, J. Debus, and T. Bortfeld. A new concept for interactive radiotherapy planning with multicriteria optimization: First clinical evaluation. *Radiotherapy and Oncology*, 85(2):292–298, 2007.
- [57] K. Todi, D. Weir, and A. Oulasvirta. Sketchplore: Sketch and explore with a layout optimiser. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, pp. 543–555, 2016.
- [58] E. Tufte and P. Graves-Morris. The visual display of quantitative information.; 1983, 2014.
- [59] J. J. Van Wijk. Views on visualization. *IEEE transactions on visualization and computer graphics*, 12(4):421–432, 2006.