



HAL
open science

A Unifying Computational Model of Decision Making

Alexandra Kirsch

► **To cite this version:**

| Alexandra Kirsch. A Unifying Computational Model of Decision Making. 2019. hal-01693687v3

HAL Id: hal-01693687

<https://hal.science/hal-01693687v3>

Preprint submitted on 4 Jan 2019

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

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

A Unifying Computational Model of Decision Making

Alexandra Kirsch

Received: date / Accepted: date

Abstract Decision-making has long been of interest as a descriptive phenomenon in psychology and as a generative one in artificial intelligence. Research ranges from general, descriptive models of heuristic decision-making to detailed studies of decision parameters. This paper introduces a model that formalizes and integrates several descriptive models so that it can both serve as a framework for psychological models and as an algorithm for computational decision-making. We set special focus on the instantiation of this model with respect to aggregating cue values by reviewing some methods from the field of computational social choice. To show its applicability in the context of artificial intelligence we present a case study of computational problem solving.

1 Introduction

The human capacity of making appropriate, timely decisions in changing, partially unknown environments has been fascinating cognitive scientists and researchers in artificial intelligence alike. Heuristics seem to be the key to this unique capacity. Despite occasional flaws (Ariely, 2010), heuristics generally lead to ecologically valid decisions (Gigerenzer and Brighton, 2009).

Starting with the seminal work of Tversky and Kahneman (1974), researchers soon found a plethora of heuristics people use when making decisions. Soon the need was felt to bring order into the ever growing list of heuristics, finding commonalities and differences, as well as to develop a more thorough understanding (Figure 1). Svenson (1979) classifies decision rules along several dimensions. Gigerenzer (2001) has proposed the “Adap-

tive Toolbox”, characterizing heuristics along three modules: search rules, stopping rules and decision rules. Shah and Oppenheimer (2008) explain heuristics in an effort-reduction framework, where a basic approach to decision making is analyzed with respect to cognitive costs and how heuristics reduce such costs.

All of these attempts are based on empirical evidence of single heuristics; all of them are descriptive in style. In this article, we incorporate these models into a modular computational model. We show how our model includes the previous models, thereby establishing its explanatory power of human behavior. The model behavior is determined by a set of parameters. Depending on the setting of these parameters, the model can be instantiated to the original heuristics that have been shown to be used by humans.

A computational model is not only more precise, it also provides an implemented decision procedure to engineers that may help to make more human-like decisions and thus improve human-machine interaction. The engineering aspect is related to the tradition of cognitive architectures (Langley, 2017; Kotseruba and Tsotsos, 2016), trying to model human intelligence as well as using the models in artificial intelligence applications.

Along with the overarching descriptive models, ongoing research has attempted to determine the choice of different strategies or to model underlying cognitive processes (Marewski and Mehlhorn, 2011). Such approaches differ in their assumptions of the underlying representation of decisions (Lee and Newell, 2011; Glöckner et al., 2014), whether strategies are distinct or emerge from parameters of the underlying decision mechanism (Rieskamp and Otto, 2006; Marewski and Schooler, 2011; Glöckner et al., 2014), or whether they are fixed or learned (Rieskamp and Otto, 2006). Most




Fig. 1: Evolution of research on single heuristics to descriptive frameworks and the generalized computational model described here. From the general model all the basic heuristics can be instantiated with appropriate parameter settings. Example heuristics taken from Gigerenzer and Brighton (2009).

of this research limits its attention to compensatory vs. noncompensatory mechanisms to choose among alternatives. Our model allows for any aggregation mechanism in the context of a full, iterative decision procedure.

While some instantiations of the generalized model correspond to heuristics that have been observed in human decision making, others may be just useful for engineering purposes. To go beyond the compensatory vs. noncompensatory debate in human decision making, we review voting methods from the field of computational social choice (Section 3), which aggregate ordinal-scale cues, in addition to the well-known summation of numeric cues. These methods have been developed in a different context without any intention to model human behavior. We will see that some of these methods happen to correspond to heuristics from human decision making. Others may be interesting from a purely computational point of view. And maybe some provide new impulses for research on human decision making.

Our model provides a generalized view on heuristics and an implemented computational framework. The parameters, however, still have to be found for each application in turn. We use the Traveling Salesperson Problem as an example to demonstrate the versatility of our model in a multi-step decision procedure. The chosen parameters are not necessarily the best for modeling the Traveling Salesperson Problem, but they are good enough to show how different parameter settings affect the outcome and how the behavior follows heuristics from the literature.

Finally, we present a generalization of the model that makes any parameter adjustable in every decision step. By this transformation, the choice of parameters (which corresponds to the use of a specific heuristic) becomes simply another decision task, and can therefore be decided with the mechanism of the model.

2 The Computational Model

Heuristics are usually described as distinct strategies for decision making. Some frameworks, as mentioned in the introduction, have tried to classify heuristics and explain heuristic decision-making by basic principles that are adapted according to the task, the situation and the decision maker. We first transform two approaches from the literature into a more formalized form. Then we derive a general computational decision procedure that incorporates both approaches and discuss how its components form different heuristics. In a next step we extend the model to allow for flexible choice of strategies and learning. We then discuss the model with respect to the adaptive toolbox (Gigerenzer, 2001) and Svenson (1979)’s 1979 classification.

2.1 Formalization of Heuristics

A decision consists of choosing an *alternative* a^* from a set of alternatives $a_i \in A$ (which may be given or may be retrieved by the decision-maker). The quality of each alternative is assessed by a number of *cues* $c_i \in C$. We formalize cues as functions in two ways: 1) as a *cue value function* mapping alternatives to real values, $v : A \rightarrow [0.0, \dots 1.0]$. The restriction to the interval between 0 and 1 is just a convenience to make sure that cues map to a comparable range, otherwise the range of a cue could introduce an implicit weight. 2) as a *cue ranking function* returning a preference ranking, which is a linear ordering \succeq of A .

Shah and Oppenheimer (2008) build their effort-reduction framework on the weighted additive rule and divide the decision process into five tasks:

0. A vector of alternatives is given
 $\mathbf{a} = (a_1, a_2, \dots, a_n), a \in A^1$

¹ (Shah and Oppenheimer, 2008) do not explicitly mention this step, we added it as step 0 for consistency with our generalized model.

1. “Identifying all cues—all relevant pieces of information must be acknowledged.” (Shah and Oppenheimer, 2008, p. 1, item 1). We formalize this with a function GET-CUES that returns a vector of cues (i.e., value or ranking functions) $\mathbf{c} = (c_1, c_2, \dots, c_m)$
2. “Recalling and storing cue values—the values for the pieces of information must either be recalled from memory or processed from an external source.” (Shah and Oppenheimer, 2008, p. 1, item 2). The cue values in this quote correspond to the result of a cue value function, thus, for the weighted additive rule we restrict cues to cue value functions. Applying the cues to all alternatives is equivalent to filling a decision matrix M :

$$\begin{array}{c}
 \begin{array}{cccc}
 & a_1 & a_2 & \dots & a_n \\
 c_1 & c_1(a_1) & c_1(a_2) & \dots & c_1(a_n) \\
 c_2 & c_2(a_1) & c_2(a_2) & \dots & c_2(a_n) \\
 \dots & \dots & \dots & \dots & \dots \\
 c_m & c_m(a_1) & c_m(a_2) & \dots & c_m(a_n)
 \end{array}
 \end{array}$$

3. “Assessing the weights of each cue—the importance of each piece of information must be determined.” (Shah and Oppenheimer, 2008, p. 1, item 3). This is another retrieval process. We treat weights as a special case of parameters that may be necessary in the process (such as parameters of the cue functions). Therefore, we call the retrieval function GET-PARAMS, which returns (possibly among other things) a vector of weights \mathbf{w} , $|\mathbf{w}| = m$.
4. “Integrating information for all alternatives—the weighted cue values must be summed to yield an overall value or utility for the alternative.” (Shah and Oppenheimer, 2008, p. 1, item 4). This corresponds to a multiplication of the filled decision matrix with the weight vector, resulting in a vector of utilities $\mathbf{u} = M^T \cdot \mathbf{w}$, $|\mathbf{u}| = n$.
5. “All alternatives must be compared, and then the alternative with the highest value should be selected.” (Shah and Oppenheimer, 2008, p. 1, item 5). This is the mathematical operation of choosing the alternative with the highest utility value $a^* = \max_{\mathbf{u}} \mathbf{a}$.

The weighted additive rule is also the standard approach in computational decision making in all types of search and optimization tasks. By making the above steps explicit, Shah and Oppenheimer (2008) point out ways in which heuristics reduce the cognitive effort, such as examining only a subset of available cues or omitting the weights.

Other simplifications require a modification of the process to allow for several iterations of decisions. The

```

def DECIDE (a, c, p):
  M ← FILL-DECISION-MATRIX(a, c)
  r ← AGGREGATE-AND-ORDER(M, p)
  a* ← FIRST(r)
  if ACCEPTABLE (a*, r, M, p):
    return a*
  else:
    DECIDE (UPDATE-ALTERNATIVES (a, M),
           UPDATE-CUES (c, M),
           UPDATE-PARAMS (p, M))

```

Fig. 2: General model of heuristic decision-making. The functions and parameters are explained in the text.

QuickEst heuristic (Hertwig et al., 1999) is a good example of such an approach:

1. A vector of alternatives is given or retrieved from memory or the environment $\mathbf{a} = (a_1, a_2, \dots, a_n)$.
2. Cues are ranked according to their “coarseness” (Hertwig et al., 1999, p. 223), which means those cues that can eliminate many alternatives quickly are used first. We thus get the cue vector $\mathbf{c}_\succ = (c_1, c_2, \dots, c_m)$ with ordered cues. In the QuickEst heuristic, cues are assumed to be boolean functions, thus mapping any alternative to the value 0 (false) or 1 (true), meaning that an alternative either has a certain property or not.
3. Repeat, starting with $i \leftarrow 1$ and $\tilde{a} \leftarrow \mathbf{a}$:
 - (a) Apply cue c_i to alternatives in \tilde{a} , resulting in a set of alternatives with positive evaluation \tilde{a}^+ and alternatives with negative evaluation \tilde{a}^- .
 - (b) If $|\tilde{a}^+| = 1$, return the only element in \tilde{a}^+ , otherwise repeat with $i \leftarrow i + 1$ and $\tilde{a} \leftarrow \tilde{a}^+$

The stopping criterion in step 3 is incomplete. One can define different ways of dealing with situations where no single alternative fulfills all available cues. QuickEst is just one example of many heuristics that decide in iterations, changing the cues and/or the alternatives in each iteration.

2.2 The Integrated Model

We combine the concepts from the weighted additive rule with iterative procedures into one algorithm, shown in Figure 2. The initial call is

```

DECIDE (GET-ALTERNATIVES (), GET-CUES (),
       GET-PARAMS ())

```

The GET-... functions are special cases of the UPDATE-... functions without initial values. These functions retrieve alternatives, cues and parameters (such as weight vectors) from memory or from the environment, the

GET- functions without any prior knowledge (except for the situation and task context), the UPDATE- functions with knowledge about the values of the last iteration.

Thus, DECIDE starts with an initial vector of alternatives \mathbf{a} , an initial vector of cues \mathbf{c} and an initial set of parameters \mathbf{p} . Let us assume that alternatives and cues contain at least one element, whereas \mathbf{p} can be empty, depending on the configuration of the algorithm. The function FILL-DECISION-MATRIX calls all cues in \mathbf{c} on all alternatives in \mathbf{a} . The cue functions may be value or ranking functions (more on this in the next section).

The function AGGREGATE-AND-ORDER combines the judgements of different cues and different alternatives and sorts alternatives so that the collectively highest ranked alternative is first. AGGREGATE-AND-ORDER may need a weight vector or other parameters, depending on the used aggregation function, as discussed in the following section. The sorting is not strictly necessary and seems very unlikely as a model of human thinking, it just facilitates the next step, which is just to take the first alternative in the ranked vector \mathbf{r} . Without the sorting, the function FIRST would have to be replaced with a function that retrieves the best alternative from the set. The typically low number of alternatives in human-like decision procedures leads only to a slight computational overhead.

The ordering step can contain many special cases that are not explicitly handled in the algorithm in Figure 2, but can be set according to the heuristic to be modeled or used. If there is only one cue in \mathbf{c} , the step contains only the sorting of alternatives according to the cue values or the alternatives are already sorted by the cue rating function. If there is just one alternative in \mathbf{a} , the step is trivial. The preceding step of filling the matrix, however, should not be omitted, because the decision of whether to accept the alternative is still to come. If several alternatives share the first rank, they will be sorted in some way in the ranked vector \mathbf{r} , maybe using a random order for equal evaluations. Again, the next step may decide to run another iteration to make a more informed choice.

The function ACCEPTABLE decides whether the best alternative is good enough to be the result of the decision process. This function can base its decision on the ranking \mathbf{r} or the pure cue values from the decision matrix \mathbf{M} , and can be configured by parameters in \mathbf{p} . Thus, an alternative may be judged unacceptable when the next-best alternative in the ranking has an equal or very close utility score. Or it may be rejected as being rated too low on certain cue values. For practical reasons, it is advisable to also include a counter in ACCEPTABLE to set a maximum number of iterations, or lower the limits for acceptance, otherwise the decision

process could run indefinitely. On top of these considerations is the question of how difficult or simple it would be to generate more cues or alternatives. With a high effort to get more cues, one might be more inclined to guess between equally ranked alternatives.

If ACCEPTABLE is not satisfied with the best alternative, DECIDE is called again, but the input parameters may be changed. The vector of alternatives might be diminished by alternatives that failed on some cue criterion (as in the QuickEst heuristic) or that fall below a certain threshold (as in the elimination by aspects heuristic (Svenson, 1979)). Alternatives may also be added if none of the previous alternatives was good enough. This could mean additional effort of memory retrieval or perception. In the same way cues can be reduced or added to. Several heuristics such as QuickEst or Take-the-best (Gigerenzer and Goldstein, 1999) use one cue per iteration, thus UPDATE-CUES would return the next cue. One may re-consider cues (see Rieskamp and Hoffrage (1999, p. 150)) that have been used in previous iterations. The weight vector is usually assumed to be static, but the algorithm leaves the possibility open to also adapt weights or other parameters.

2.3 Discussion of the Model

We motivated the model by formalizing the basic decision procedure by Shah and Oppenheimer (2008) and the QuickEst heuristic (Hertwig et al., 1999) as an instance of iterative heuristics, both of which were derived from empirical evidence of human decision making. We have constructed the generalized model to include them. Therefore, the generalized model has some empirical backup, at least for those parameter settings that instantiate the respective framework and heuristic. Let us first make sure that the algorithm in Figure 2 really comprises the two.

For the Shah and Oppenheimer (2008) model, a list of fixed alternatives can be provided with the function GET-ALTERNATIVES (step 0); GET-CUES corresponds to the retrieval of cues (step 1), GET-PARAMS models the retrieval of the weight vector (step 3). The retrieval and storage of cue values (step 2) corresponds to FILL-DECISION-MATRIX, and their aggregation (step 4) as a weighted sum is one way of instantiating AGGREGATE-AND-ORDER. The choice of the best alternative (step 5) is simplified by the ordering included in AGGREGATE-AND-ORDER. The Shah and Oppenheimer (2008) model assumes only one decision cycle. Therefore, ACCEPTABLE will accept the alternative with the highest aggregated cue value, the UPDATE- functions are irrelevant.

For QuickEst, step 1 again corresponds to GET-ALTERNATIVES being instantiated to returning a

list of fixed alternatives. For GET-/UPDATE-CUES we need a vector of ordered cues c_{\succ} . The functions GET-/UPDATE-CUES return one element per round in the order given in c_{\succ} . FILL-DECISION-MATRIX is reduced to applying the one cue chosen for the iteration to all the alternatives, filling the matrix with boolean values (or integers 0 and 1). The function ACCEPTABLE accepts a choice when only one alternative comes out true, otherwise it requires another iteration. UPDATE-ALTERNATIVES returns all alternatives that were rated as true in the last iteration.

Figure 1 lists two other frameworks that generalize observed heuristics. One is the Adaptive Toolbox by Gigerenzer, which follows a similar idea of generalizing heuristics into a unified method and depending on the parameters instantiates to different heuristics. The Adaptive Toolbox can be mapped to our model as follows:

Search rules Gigerenzer (2001) distinguishes the search for alternatives, as modeled in the GET-/UPDATE-ALTERNATIVES functions, and the search for cues, as in the GET-/UPDATE-CUES functions. The options he lists (“random search, ordered search (e.g., looking up cues according to their validities), and search by imitation of conspecifics, such as stimulus enhancement, response facilitation, and priming” (Gigerenzer, 2001, p. 44)) would have to be implemented by different versions of these functions.

Stopping rules are implemented by the function ACCEPTABLE. They can, for instance, be based on absolute aspiration levels as in satisficing models (Simon, 1956) or if a cue is found that favors one alternative as in QuickEst or other one-good-reason heuristics (Gigerenzer and Goldstein, 1999).

Decision rules correspond to AGGREGATE-AND-ORDER, i.e. the aggregation of cue values. Gigerenzer (2001) describes them after the stopping rules, whereas our model performs this step before calling ACCEPTABLE. If ACCEPTABLE is implemented in a way that it can decide without knowing the aggregated values, the algorithm may be modified to first test for acceptability and then call AGGREGATE-AND-ORDER only in the affirmative case. But sometimes one may base this choice on aggregated values, as in our case study below, where we use an acceptability criterion that compares the aggregated value of the best choice to the aggregated value of the second best choice. Thus, the algorithm in Figure 2 is slightly more general, but also less efficient in cases where ACCEPTABLE works with unaggregated values.

The third framework, or rather classification, from Figure 1 is the one by Svenson (1979). It classifies decision rules along three dimensions, which correspond to our model in the following ways:

Metric level of cues: We have differentiated between cue ranking and cue value functions, where ranking corresponds to an ordinal scale and values to ratio or interval scale. The ordinal scale, however, is not treated as a mere ordering by Svenson (1979), some rules use criterion values corresponding to absolute thresholds. The instantiation of AGGREGATE-AND-ORDER must be consistent with the output of the cue functions. Different ways of combining ordinal values are presented in the next section.

Lexicographic order: This dimension says whether cues are used all at once (which would mean that no order is necessary) or whether they are used in an iterative fashion, one per iteration, in which case the cues must be ordered. We have shown in the QuickEst heuristic, how this can be implemented with our model.

Commensurability: This dimension classifies rules according to whether a high score by one cue can counterbalance a low score by another. Weighted sums are an example of commensurable aggregation, whereas QuickEst is non-commensurable (because once an alternative has been eliminated it will never be the final choice, no matter how well it would be rated by subsequent cues). This quality is determined by an interplay of AGGREGATE-AND-ORDER, ACCEPTABLE and the control of the iterations by UPDATE-ALTERNATIVES and UPDATE-CUES.

Svenson (1979) points out that decision rules could change or be adapted during the decision process. Section 5 presents a generalized version of our model that includes this functionality.

Just like the other frameworks, our model can be instantiated to specific heuristics. Table 1 lists three example heuristics from the literature and how they are instantiated in our model. Such an analysis opens new views on well-established heuristics. For example, Take-the-best and Tallying are described with binary cue values, but could easily be generalized to numeric cue values.

The literature discusses other options for decision making than heuristics. Figure 3 classifies the broad lines of decision making theories. In the following we show how rational decision making and decisions by recognition are encompassed in our model.

Rational decision making is another way of choosing among alternatives. It assumes that the decision maker has access to all possible alternatives and should ideally choose the one with the highest expected utility.

FILL-DECISION-MATRIX	AGGREGATE-AND-ORDER	ACCEPTABLE	UPDATE-ALTERNATIVES	UPDATE-CUES
Take-the-best (Gigerenzer and Brighton, 2009)				
“To infer which of two alternatives has the higher value: (a) search through cues in order of validity, (b) stop search as soon as a cue discriminates, and (c) choose the alternative this cue favors.”				
(binary) cue value function for one cue c_i	SORT-DESC (M) [M has only one row]	$c_i(a^*) > c_i(\text{SECOND}(\mathbf{r}))$	– [alternatives given]	get next cue from a list of cues ordered by validity
Tallying (Gigerenzer and Brighton, 2009)				
“To estimate a criterion, do not estimate weights but simply count the number of positive cues”				
binary cue value function	SORT-DESC ($M^T \cdot \mathbf{1}$)	True [always accept highest alternative]	–	–
Elimination by aspects (Svenson, 1979)				
“First, the most important [cue] is investigated, and all alternatives that do not exceed the conjunctive criterion on this [cue] are eliminated. This procedure is then repeated with new [cues] successively lower in the lexocographic order”				
cue value function for one cue c_i	SORT-DESC (M) [M has only one row]	$c_i(a^*) \geq T \wedge c_i(\text{SECOND}(\mathbf{r})) < T$ $T \in \mathbf{p}$ is a given threshold value	$\{a c_i(a) \geq T\}, T \in \mathbf{p}$	get next cue from a list of cues ordered by validity

Table 1: Examples of heuristics instantiated in our model. SORT-DESC denotes a function sorting the alternatives in descending order of their cue values.




Fig. 3: Ontology of decision strategies with examples of references holding the particular view.

Each alternative leads to possible outcomes r_1^a, \dots, r_m^a with probabilities p_1^a, \dots, p_m^a , and each outcome has a subjective value of desirability $\nu(r)$. The utility of each alternative i is calculated by

$$U(a) = \sum_{k=1}^m (\nu(r_k^a) * p_k^a)$$

A rational decision maker chooses the option with maximum utility. The desirability is usually given as a weighted sum of features

$$\nu(r) = \sum_i w_i f_i(r)$$

If we multiply the weight vector \mathbf{w} with the probability distribution, we get a new weight vector that includes both desirability of a feature and the probabilities of the different outcomes. Thus, rational choice is

a heuristic choice where 1) all possible alternatives are considered and 2) the weight vector contains probabilities (if probabilities are modeled at all, otherwise it is identical to the original weight vector).

Rational choice in multi-step decision problems is implemented in optimization and search algorithms. For coping with large state spaces, such algorithms also reduce the alternatives considered, like the simplex algorithm for linear programming (Dantzig, 1963) and all types of best first search (Russell and Norvig, 2003), such as greedy and A* search. While these algorithms reduce the number of alternatives by exploiting structural properties of the state space, as defined by the objective function, our model assumes retrieval of alternatives from memory or the environment. The model can be instantiated to standard optimization, where the reduction of alternatives serves only to accelerate the process, trying to find the optimal solution or one in a close boundary to the optimum. But in the general case, the model strives for a satisficing solution, rather than an optimal one (which may not even be definable in real-world tasks).

Choice by recognition assumes that no alternatives are compared, but that the decision maker chooses the alternative that first comes to mind. This approach is a special case of a take-the-best or recognition heuristic. In our model, the decision maker would generate only one alternative. The challenge of decision making by recognition is the memory access and distance functions to decide whether situations are similar enough to

reuse a prior solution. This challenge is the same when generating more than one alternative in our decision model.

In sum, the computational model of Figure 2 integrates several descriptive frameworks and categorizations of heuristics into one. The heuristics that underlie the frameworks and therefore the resulting model, are based on empirical evidence of human decision making. The single heuristics can be reproduced by the model with specific sets of parameters. Not all parameter choices, however, lead to heuristics that are employed by people. Some choices may be useless, others may lead to good computational performance without imitating human behavior.

3 Aggregation of Cues

In this section we take a closer look at the representation and aggregation of cues. The discussions on human decision making mostly focus on compensatory vs. noncompensatory methods. Also since many experiments restrict choices to two alternatives, this discussion has been rather restricted. Here we offer a wider viewpoint from another scientific field: computational social choice. These methods make no claims of being used in human decision making, but they may provide impulses for ongoing discussions on human decision making as well as opening possibilities for engineering uses of the model.

The classical approach to aggregate the judgements of all cues to determine the overall best alternative is a weighted sum, or the simpler variant, an unweighted sum (which corresponds to a weighted sum when all weights have the same value). But still we would need all cues to be implemented as a cue value function, which means they have to provide numeric values. In many cases it seems simpler to just rank the alternatives by preference without giving them explicit values. But how can we aggregate rankings?

The field of computational social choice (Brandt et al., 2016) is concerned with developing fair voting systems. The assumption there is that each voter from a finite set of voters² V , $|V| = m$ is given a finite set of alternatives A , $|A| = n$. Each voter i casts a *ballot* (or ranking), which is a linear ordering \succeq_i of A . A *profile* $P = (\succeq_1, \succeq_2, \dots, \succeq_m)$ specifies a ballot for all voters in V ; $\mathcal{L}(A)^m$ denotes the set of all such profiles for a given m . A *social choice function* maps a profile to a single combined preference ranking $\text{scf} : \mathcal{L}(A)^m \rightarrow \mathcal{L}(A)$.

² The notation and definitions are taken from Zwicker (2016).

If we now replace the set of voters V by the set of cues C in our decision procedure, we can use social choice functions to combine cue rankings into an overall ranking of the alternatives. However, there is a fundamental difference in the tasks of voting and cue aggregation: voting assumes equality of voters (each has one vote), whereas cues are often assumed to be of unequal importance. For example, Russo and Doshier (1983) assume that the least important cue is “the one with the smallest dimensional difference” (Russo and Doshier, 1983, p. 683). Huber (1979) mentions the relation of his weighted-sets-of-dimensions model to social choice functions and suggests an interpretation of \succeq as a “partial order of social power” (Huber, 1979, p. 163). Even though social choice is not motivated by cognitive mechanisms, there are certain connections. For example, Katsikopoulos and Martignon (2006) draw an explicit connection between heuristics and Condorcet theorems.

Cue value functions can be transformed into cue ranking functions by sorting the alternatives for each cue according to the value assigned by the cue. Thus, a value aggregation function (i.e. summing) can be used as a social choice function, by sorting the alternatives according to the combined score. The reverse case (from ranking to value) is possible, but under-specified. Scoring functions offer different ways of assigning cue values to ranked alternatives.

Scoring Scoring rules assign a value to each alternative in a ranking, thus transforming a ranking function into a value function. The methods differ in the way they assign the numbers, here are two examples:

Borda count: $f(r) = n - r$, where r is the rank of the alternative and n is the number of all alternatives. Thus, the highest ranked alternative receives value $n - 1$, the second highest $n - 2$, the lowest one 0.

Formula One Championship: The values 25, 18, 15, 12, 10, 8, 6, 4, 2, 1, 0, ... 0 are assigned to the ranked alternatives in the given order, i.e. the highest ranked alternative is assigned the value 25, the second highest 18, etc.

For a comprehensive list of scoring methods we refer to Zwicker (2016, p. 36 ff.). The transformed cue values are aggregated with an unweighted sum.

Voting Voting methods combine rankings directly without the intermediate step of assigning scores. For two alternatives and odd number of cues, majority rule is the uncontroversial choice (Brandt et al., 2016, p. 34, proposition 2.2).

For more than two alternatives, most voting rules are based on the concept of *net preference*. The net preference of alternative a_1 over alternative a_2 is defined as

$$Net(a_1 > a_2) = |\{c \in C | a_1 \succ_c a_2\}| - |\{c \in C | a_2 \succ_c a_1\}|$$

i.e. the number of cues that rank a_1 over a_2 diminished by the number of cues that rank a_2 over a_1

Note that even though the decision matrix is filled row-wise (cue-wise), the net preference compares alternatives, resulting in a matrix of preferences for alternatives:

$$\begin{array}{c} \begin{array}{cccc} & a_1 & a_2 & \dots & a_n \\ a_1 & - & Net(a_2 > a_1) & & Net(a_n > a_1) \\ a_2 & Net(a_1 > a_2) & - & & Net(a_n > a_2) \\ \dots & & & - & \\ a_n & Net(a_1 > a_n) & Net(a_2 > a_n) & & - \end{array} \end{array}$$

An alternative a that defeats every other alternative in the strict pairwise majority sense $a > a_i$ for all $a_i \neq a$ is called a *Condorcet winner*. It is generally assumed that a Condorcet winner is a fair choice in an election. However, not every profile has a Condorcet winner, but it seems desirable that a social choice function returns the Condorcet winner if one exists. Such a function is called a *Condorcet extension*. The Borda and Formula-One-Championship scoring methods are not Condorcet extensions. Here are two examples of voting rules that are:

Copeland method: Compute the Copeland score for each alternative: $Copeland(a_i) = |\{a_j \in A | Net(a_i > a_j)\}| - |\{a_j \in A | Net(a_j > a_i)\}|$, i.e. count all alternatives that are worse than a_i and subtract the number of alternatives that are better than a_i according to the net preference. Select the alternative with the highest Copeland score.

Sequential Majority Comparison is a pairwise comparison of alternatives:

1. pick some order of alternatives: (a_1, a_2, \dots, a_n)
2. $winner \leftarrow a_1$
3. **for** $i = 2 \dots n$:
 if $Net(a_i > winner)$ **then** $winner \leftarrow a_i$
4. **return** $winner$

We treat Sequential Majority Comparison as an atomic social choice function. In the decision model, one could treat it as an iterative method, in which alternatives are excluded in each iteration. However, Sequential Majority Comparison assumes that all alternatives are compared sooner or later so that the iterative approach does not give any benefit here.

More on voting rules and Condorcet extensions can be found in Zwicker (2016, p. 33 ff.).

Multiround rules Another class of voting rules compares alternatives in rounds, removing alternatives in each round and re-ranking the remaining alternatives. This approach is similar to the elimination by aspects heuristic (Svenson, 1979, p. 90), which we use below. For example, the Nanson rule (a Condorcet extension) eliminates in each round the alternatives with a Borda score below the average Borda score. For more details on multiround rules refer to Zwicker (2016, p. 37).

4 Case Study

The goal of our computational model is on the one hand to integrate existing models of heuristic decision-making, as we have already discussed. On the other hand, we want it to be usable as a component in computational decision-making systems. In this section we present experiments for a specific task, the Traveling Salesperson Problem. Rather than trying to find the one and only parameter set for this problem, we use this example to show how different parameters affect the result and how they resemble heuristics from human decision making. As another example, we have used the model for human-aware robot navigation (Kirsch, 2016, 2017).

We first introduce the task of the Traveling Salesperson Problem, then we address some implementation details, and finally show experiments with different parameter combinations of the algorithm.

4.1 The Traveling Salesperson Problem

The Euclidean Traveling Salesperson Problem (TSP) is the task of finding the shortest possible tour through a set of given 2D points, and returning to the starting point. Apart from being a representative of the class of NP complete problems (i.e. problems where the number of potential solutions grows exponentially with the size of the problem instance), it has been an object of research on human problem solving (MacGregor and Chu, 2011). Optimization programs can nowadays solve instances with thousands of points optimally and problems with millions of points near-optimally. These approaches usually tackle the whole task at once, whereas people start with one point and add connections one after the other (Tenbrink and Wiener, 2009) and find surprisingly good solutions (MacGregor and Ormerod, 1996). For related tasks that involve user interaction, a human-like approach seems desirable. Therefore, we treat a TSP solution process as a sequence of decisions of which point to add next to the tour.




Fig. 4: Representation of TSP instances and states in the solution process. The numbers next to the points are labels. The partial solution starts at point 1 and includes points 6 and 5. Point 5 is the current point from which to decide the next step. Points 4 and 7 are alternatives for the next decision.

One characteristic of TSPs is that the instances differ a lot, so that it is hard to find one parameterization that solves any instance equally well (Kirsch, 2011, 2012). This also puts some limit on the optimization of parameters, because a set of parameters might work well for one TSP instance and bad for another.

For the experiments we use an open dataset³ that contains over 25,000 human solutions to 90 TSP problems (Rach and Kirsch, 2016).

4.2 Implementation

We implemented the general decision algorithm in a system called Heuristic Problem Solver (HPS) (Kirsch, 2016). The source code used for the experiments below is available at <https://bitbucket.org/kirschalexandra/heuristicproblemsolver/commits/tag/compSoC-TSP>.

A TSP instance is represented by a set of 2D points. A state in the solution process is a partial solution, which is initially empty and contains all the points plus one (the starting point has to be the same as the end point) in the end. The last point in a partial solution is the current point, from which the next step of the tour has to be decided (Figure 4).

The implementation of `GET-ALTERNATIVES` is constant in all experiments. In HPS alternatives are generated by so-called *producers*, functions that return alternatives according to the state of the world. We used three producers that decide, which point should be the start point (Rach and Kirsch, 2016): one that returns

the three closest unvisited points to the current point, one that suggests the next point on the convex hull (the outline of the problem), and one that suggests the starting point to return to when the rest of the tour has been constructed.

The cues are represented in HPS by so-called *evaluators*, functions that receive all alternatives in the iteration and return an assignment of alternatives to real values between 0 and 1. We have a total of 10 evaluators available (the abbreviations are used in Figure 8):

nearest-neighbour (nn) prefers alternatives that are close to the current point.

last-chance (lc) tries to avoid that a point skips its nearest neighbour. Sometimes points are “left out” on the way and need to be collected in the last steps, which inevitably leads to long tours with crossings.

remaining-acc-dist (rd) estimates the accumulated distance to the remaining points and chooses the one with the lowest estimated future path length. This is a similar rationale as the heuristic function in an A* search.

no-intersections (ni) rejects alternatives that would lead to an intersection by adding the next edge.

avoid-splitting (as) punishes points that would lead to a line through the middle of the problem, thus splitting it into two halves that could only be combined later by a crossing line.

start-intersection (si) avoids alternatives whose direct connection to the starting point would cause an intersection.

follow-lines (fl) prefers alternatives that lead to a wide angle of the edges at the current point, rather than alternatives that lead to sharp turns.

convex-hull (ch) prefers moving along the convex hull.

regions (rg) uses a clustering of points into regions. In the TSP literature, hierarchical approaches have long been discussed (Graham et al., 2000). This evaluator prefers alternatives that are in the same region as the current point.

regions-convex-hull (rc) like *regions* prefers points in the same region, but also points in the region that follows the current one on the convex hull.

Not all cues are equally useful, and it may depend on the aggregation of cues, which ones should be used. Therefore, in the experiment we ran an optimization for each cue aggregation function to find the best set of cues, and for weighted sums the weight for each cue. In the first experiment, the `GET-CUES` function returns a fixed set of evaluators, depending on the used aggregation function. `GET-PARAMS` returns an association of cues to weights, if necessary.

³ <http://www.wsi.uni-tuebingen.de/lehrstuehle/human-computer-interaction/home/code-datasets/tsp-dataset/perlentaucher-2.html>

In the first set of experiments, the functions UPDATE-ALTERNATIVES and UPDATE-CUES are irrelevant, because only one iteration is run. For the second set of experiments, the functionality of these functions is explained below.

Different options for AGGREGATE-AND-ORDER are tested in the first set of experiments. ACCEPTABLE is by default a function that accepts any alternative. In the second experiments we use another definition that allows for an iterative process.

4.3 Aggregation Functions

First, we test the rather classical approach of evaluating a given number of alternatives with a given number of cues, thus filling the whole decision matrix once. We look at the effect of different aggregation functions to determine the best choice:

summing: *weighted sum, unweighted sum* (weighted sum with equal weights)

ranking: *copeland, sequential majority comparison*

scoring: *borda, formula-one championship*

All of these six strategies have parameters: weighted sums assign a weight value between 0 and 1 to each cue, for simplicity we assume steps of 0.1, which results in a parameter space with a size in the order⁴ of 10^{10} . For the other methods, there is only the choice of whether to use a cue or not, thus each cue is assigned a Boolean value, resulting in a parameter space of $2^{10} - 1 = 1023$ (the empty set of cues is excluded).

Parameter Optimization From the 90 problems in our TSP data set, 20 are pairs of problems with identical geometrical layout, but slight variations in presentation (e.g. one task is presented with differently colored points, the other with one color for all points). From these pairs, one version per pair was used as a training set, thus we had 10 training tasks for the optimization process. The remaining 70 tasks served as test problems.

The solution quality of the Traveling Salesperson Problem is usually measured by the percentage above the optimal tour length (PAO). For each TSP instance, only a finite number of specific PAO values is possible (since there is a finite number of tours). Therefore, the PAO values of different TSP instances are incomparable and should not be aggregated. To allow for some aggregation of the 70 PAOs for each tour, we normalize the PAO value among all solutions that were generated

⁴ 10^{10} is an upper bound, because some configurations are equivalent, for example when all weights have the same value.

in the following experiments with a student transformation: $p'_i = \frac{p_i - m_i}{s_i}$, where p_i is the percentage above optimum for a specific solution of TSP instance i ; p'_i is the normalized value; m_i is the mean PAO of all solutions for instance i , and s_i is the (descriptive) standard deviation of PAO values over all solutions of instance i . This normalization allows for a comparison among different instantiations of the decision procedure, but it is inappropriate and unintended for assessing the absolute solution quality.

For the five conditions that are controlled only by the number of cues used, we used the brute force method of solving each of the 10 training tasks with each of the 1023 parameter sets, normalizing over the solutions per task, aggregating the PAO values for each parameter set, and choosing the one with the lowest relative tour length. This takes about 20 minutes per condition on a desktop computer (Intel Core i5-2520M (2.50GHz), 8 GB RAM).

For the weighted sum, as the parameter space is far too large for the brute force approach, we used a genetic algorithm (Banzhaf et al., 1998) with a population size of 16 and 64 generations, resulting in 1024 tested parameter combinations, thus having an optimization effort comparable to the other conditions. The fitness of each parameter combination was again assessed by aggregating normalized PAO values.

To evaluate the robustness of the aggregation functions, we ran the same optimizations with two other measures of tour quality: 1) a rating that measures the similarity of a solution to human solutions of the same task, and 2) the count of how many participants in the data set chose the same solution for a tour. In addition, for the weighted sum we experimented with different numbers of generations of the genetic algorithms to assess the necessary effort for optimization.

Quality of Decisions Figure 5 shows the comparative results over the 70 test problems, lower values denote shorter, and therefore better, tours. The median under all conditions is below 0 (which would be the average relative performance per tour), which shows that in general the performance is on some common level, but that all conditions can produce exceptionally long tours. The performance in all conditions is practically the same when the normalized PAO is aggregated. This does not mean, however, that they all produce the same results on the same tasks. Figure 6 shows the unnormalized PAO values per task for the first 20 tasks from the test set.

Optimization Effort For this experiment we more or less arbitrarily decided on using 1024 parameter sets

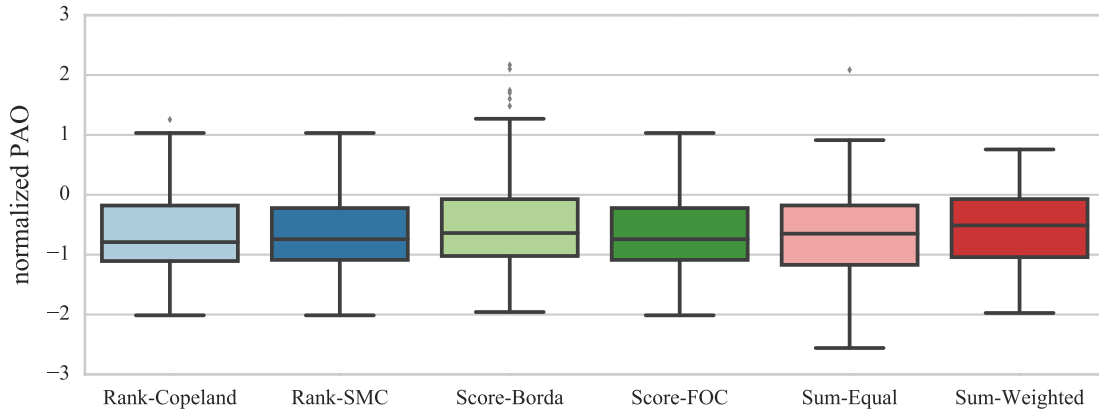


Fig. 5: Quality with optimized configurations over 70 test tasks. Lower values denote shorter, i.e. better, tours.

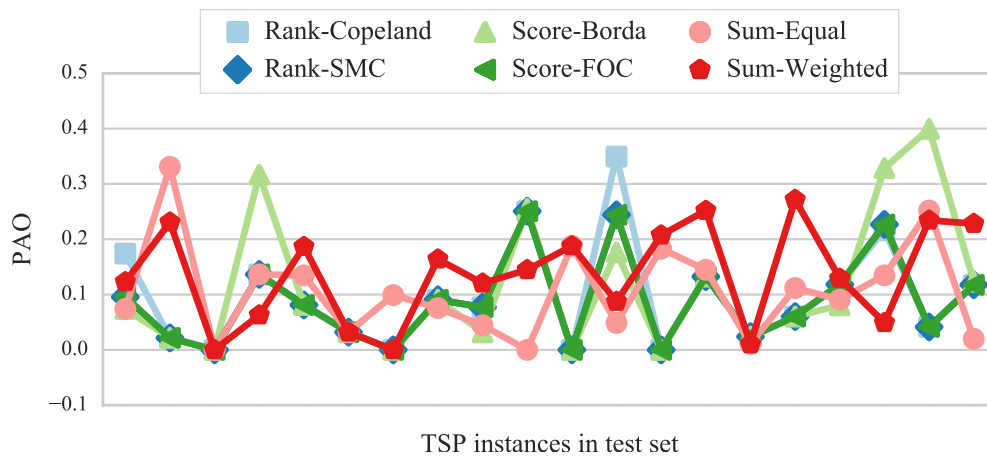


Fig. 6: (Unnormalized) PAO of optimized configurations for 20 test tasks.

tings in the optimization. Possibly, a similar level of performance could be reached with less optimization effort, or an additional effort could boost the performance. Figure 7a shows the performance for the weighted sum on the test data when trained in 16, 32, 64 and 128 generations. The performance is practically identical for all variants. To better understand the reason, Figure 7b shows the performance on the training data. This shows that the algorithm is improving on the training data with more generations, but this has no effect on the performance on the test data.

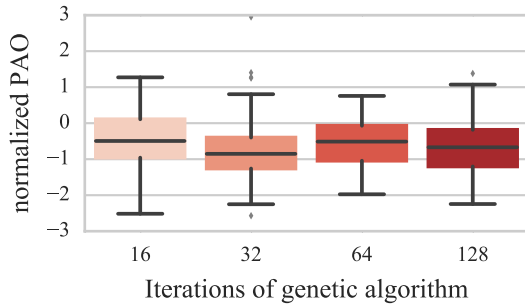
This overfitting can be explained by the high diversity of TSP instances. In such a setting, a costly optimization seems not to pay off, because instances require different parameters. This confirms the finding of others (Keller and Katsikopoulos, 2016) where in real-world tasks parameter optimization is not possible, or even necessary.

Influence of Parameter Choice Another question is how important the parameters are for making good decisions. In the optimization, all variants of aggregation functions faced the same challenge of the diversity of TSP instances. But how much is the performance influenced when the value to be optimized for is a different one than that is tested for? Or how does the performance change with a non-optimized configuration?

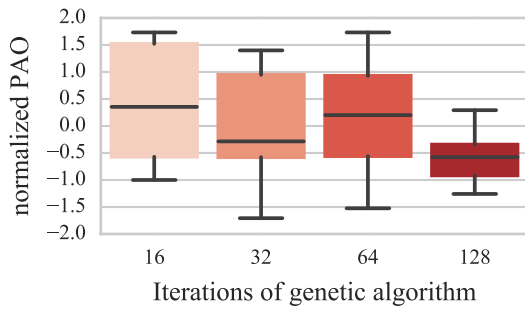
Figure 8 analyzes which cues were used in the optimized configurations. It shows that there are some cues that were used more often than others. It also shows that the weighted sum used 9 or 10 cues, no matter for which target value it was optimized, whereas the conditions with a boolean choice of cues mostly used 3–5 cues.

From these observations we constructed the following parameter sets:

- *opt pao/ rating/ count*: the individual parameter set for the aggregation function when optimized



(a) Tested on test data.




(b) Tested on training data.

Fig. 7: Results for weighted sum after different number of iterations in genetic algorithm. 64 iterations correspond to the optimization effort of the non-weighted aggregation functions.

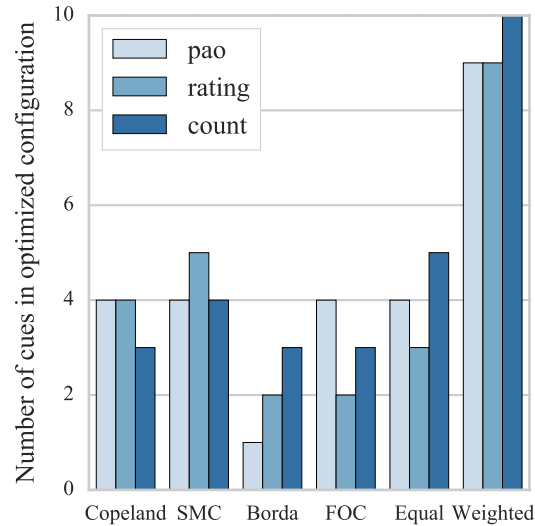
for PAO, rating of similarity to human solutions, and the count how often a solution was chosen by participants

- *four/ five/ all*: Using the four/ five cues that were used most often in the optimization process (*four*: dark blue cues in Figure 8a, *five*: dark and medium blue cues in Figure 8a); *all* uses all available cues. For the weighted sum for each condition, six randomly generated sets of weights were used; Figure 9a shows the median of those six variants, Figure 9b shows the single performances.

Figure 9a shows that the scoring methods are the most robust with respect to parameter choice. The ranking methods are mostly stable, but perform notably worse when trained on the rating value. The unweighted sum shows a similar level of stability as the scoring methods, but in most cases performs slightly worse. The most fluctuation is shown in the weighted sums aggregation function. Interestingly, it performs best (tested on PAO) when trained on the number of times a solution was chosen by participants. Figure 9b suggests



(a) Number of times expert was chosen in the 18 optimized configurations.

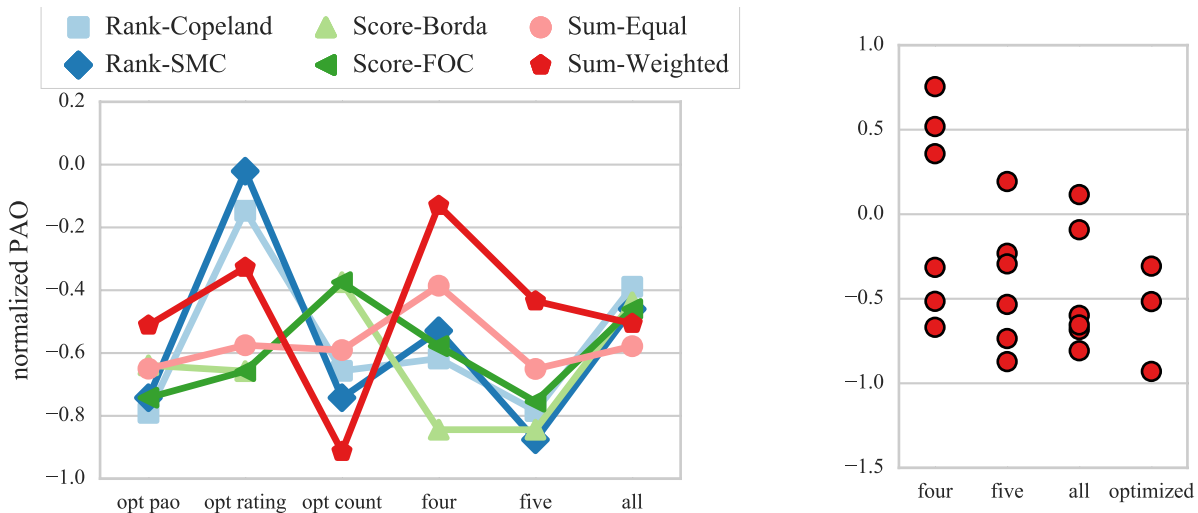


(b) Number of experts in optimized configurations.

Fig. 8: Expert usage in optimized configurations.

that this fluctuation is due to the choice of particular weights. In all three predefined parameter sets, there is a set of weights that outperforms or at least compares to the scoring methods, but there are always other choices of weights that make the performance notably worse.

Summary of Aggregation Functions The method of aggregation is mostly irrelevant in the specific task regarded here. The optimization effort we used could probably be reduced even further for all the aggregation functions. However, the weighted sum seems to need



(a) Configurations of optimization process and three constructed ones (for weighted sum: median of six random configurations).

(b) Single results of configurations with weighted sum.

Fig. 9: Robustness with respect to configuration. Lower values denote shorter tours.

slightly more attention to the parameters, the scoring methods seem to be the most stable ones. In less diverse tasks or with an adequate classification of TSP instances, the effort to learn weights for a weighted sum might pay off.

4.4 Iterative Decision Methods

We now test methods that use more flexibility of the decision model. In the last experiment, we varied the setting of AGGREGATE-AND-ORDER, now we look at some combinations of the functions ACCEPTABLE, GET-/UPDATE-CUES and UPDATE-ALTERNATIVES, considering only one cue at a time. Thus, the decision matrix has only one line per iteration, which means that AGGREGATE-AND-ORDER just orders the alternatives according to the one value provided. The other parts are configured as follows:

acceptable: We accept an alternative if $\frac{c(a^*) - c(a^+)}{c(a^+)} > \eta$, where $c(a^*)$ is the (single) cue value of the best alternative, $c(a^+)$ is the cue value of the second-best alternative, η is a parameter, which is varied in the experiment between 0.01 and 2. If there is only one alternative left or no more cues, the best (or only) alternative is chosen.

get-/update-cues: We mimic two one-good-reason heuristics (Gigerenzer and Goldstein, 1999), which use one cue per iteration:

- The *Minimalist heuristic* uses cues in random order and makes use of the recognition heuristic whenever possible (i.e. prefers the option that it

recognizes). The recognition heuristic makes no sense in TSPs, so we use the cues in random order.

- The *Take the Best* heuristic assumes that the decision maker has some knowledge about the importance of cues for a specific task and uses the cues in the order from most important to least important. Determining the “best” order of cues is similar to determining weights. We used for all tasks the order of the cues as they were included in the optimized configurations shown in Figure 8a. When several cues have the same relevance (as rd and rc), the order is chosen randomly.

update-alternatives:

- Keeping all alternatives by passing them on to the next round.
- The elimination by aspects heuristic of Svenson (1979, p. 90). In each iteration it removes low-ranked alternatives. We defined “low-ranked” as being below the mean of the cue value (unless all alternatives have the same value, then all are kept). We experimented with other parameters, such as one or two standard deviations below the mean, the results were very similar to the ones shown below.

The conditions are summarized in Table 2. GET-/UPDATE-CUES now contains some randomness (even for take-the-best as some cues have the same relevance). Therefore, each task was now solved 20 times per condition.

condition	UPDATE-CUES	UPDATE-ALTERNATIVES
<i>min</i>	minimalist	keep all
<i>elim-min</i>	minimalist	elimination by aspects
<i>tbb</i>	take-the-best	keep all
<i>elim-tbb</i>	take-the-best	elimination by aspects

Table 2: Conditions for iterative decision procedure.

Acceptance parameter The approach has several parameters, some of which are fixed and were justified in the preceding paragraph. For the parameter η of the ACCEPTABLE function, the choice is not so obvious. Figure 10 shows the performance of the *tbb* condition for different choices of η (the other conditions reveal a similar picture).

The data suggests that it is beneficial to use a low acceptance threshold. For the take-the-best update function, this makes sense as the (presumably) better cues are used first and if their decision is not used, the less reasonable cues make the decision. For the minimalist update function, this argument does not hold, but since the cues are used in random order, waiting for other cues does not in general bring any benefit either. Also from a computational point of view, few iterations are preferable.

Quality Figure 11 shows the result of testing the four conditions of Table 2 with $\eta = 0.1$, Figure 11a showing the best runs from each of the 20 runs per task, and Figure 11b the worst runs. Be aware that the y-axis shows different ranges, in Figure 11a all median values are below 0, in Figure 11b, most of them are above 0.

Condition *min* can lead both to very good and very poor results, depending on the random choice of cues. The elimination of alternatives stabilizes the behavior significantly. The poorest results of condition *elim-min* are only slightly worse than the best ones. But this stabilization also removes the positive outliers of condition *min*. Possibly, in some cases the best decision is one that is evaluated poorly by most cues and good by few. By eliminating alternatives, a kind of aggregation sets in, which removes the positive as well as the negative outliers.

Condition *tbb* performs comparably to *elim-min*. Interestingly, adding elimination of aspects leads to results similar to those of *min*. One might conclude that good results can be obtained either by a good ranking of cues (*tbb* vs. *min*) or by elimination of alternatives (*elim-min*), when such knowledge is not available. The question remains of how to obtain the positive outliers (or produce more of those) of *min* and *elim-tbb*.

Summary of Iterative Methods The parameterization of the cue ordering and whether to eliminate unpromising

alternatives, has a notable effect in the TSP domain. With the general decision procedure many more variations are possible, like using two cues per iteration, combined with different aggregation functions for those cues. Also the generation of alternatives could be explored in different directions.

We have shown some of the possibilities of using the model on the Traveling Salesperson Problem. The approaches were compared relative to each other rather than to a static baseline such as human performance data. One reason is the difficulty of finding a metric. The PAO measure that we used here is known to be unfit to compare to human solutions (Tak et al., 2008), but there is no generally agreed alternative metric for doing so. The other reason is that none of our instantiations of the model can compete with human solutions. People adapt their strategy not only to the problem, but it seems even to a particular situation in the solution process. As we will see in the next section, our model is powerful enough to encompass this change of strategies. But we know of no representation that would encode a solution step in a generalized way to allow for a situation-specific choice of parameters. The field of qualitative reasoning might offer a way to solve this in the future (Wolter and Kirsch, 2015).

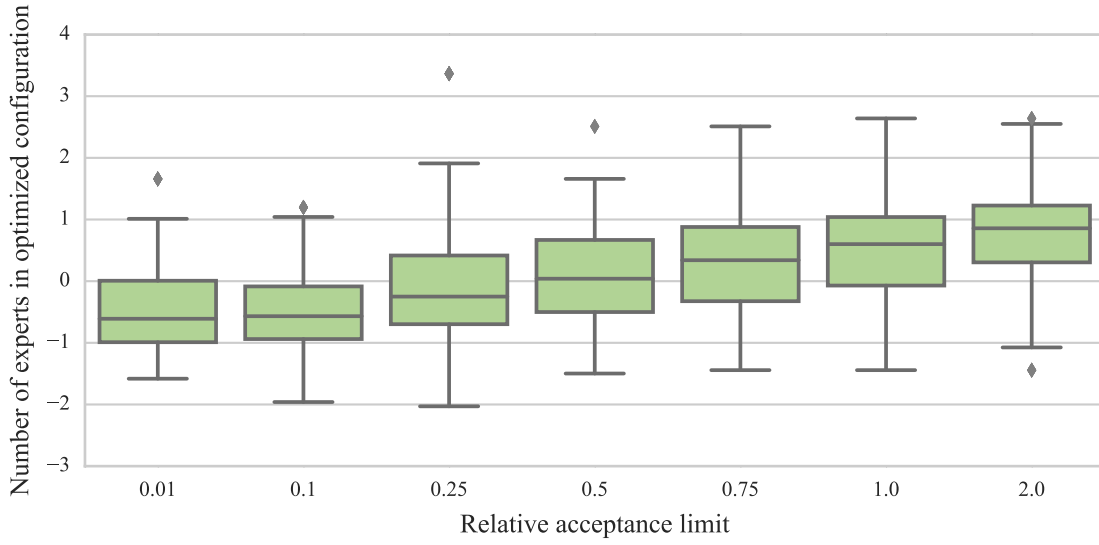

5 The Adaptive Model

The model of Section 2.2 can be parameterized into different kinds of heuristics as we have seen in the case study. However, it seems clear that people choose their strategies, i.e. heuristics, for making choices according to the situation.


We can extend our model to allow for such adaptations in each decision and even in each decision cycle. In the algorithm in Figure 12 the functions that distinguish different heuristics (AGGREGATE-AND-ORDER, ACCEPTABLE) are passed as arguments to the algorithm and can thus be changed in every iteration cycle. Especially in cognitively demanding tasks, people change their representation of the the task (Newell and Simon, 1972). Passing the UPDATE-... functions as part of the configuration, allows for such a change in the sense that strategies for retrieving cues and alternatives from memory or the environment can be adapted in each cycle.

The update of the functional parameters in this extended algorithm is performed in one function RECONFIGURE⁵. The reason is that these choices are interdependent as illustrated in Figure 13. AGGREGATE-

⁵ One could go further and also make RECONFIGURE reconfigurable.

Fig. 10: Acceptance parameter η for *ttb* condition.

(a) best runs



(b) worst runs

Fig. 11: Different variants of one-good-reason heuristics with acceptance parameter $\eta = 0.1$.

AND-ORDER together with parameters such as weight vectors determines the strategy for combining cue values. This choice has to fit the one of ACCEPTABLE (in-

```

def DECIDE (a, c, p,
            AGGREGATE-AND-ORDER,
            ACCEPTABLE,
            UPDATE-ALTERNATIVES,
            UPDATE-CUES,
            UPDATE-PARAMETERS):
    M ← FILL-DECISION-MATRIX(a, c)
    r ← CALL(AGGREGATE-AND-ORDER, M, p)
    a* ← FIRST(r)
    if CALL (ACCEPTABLE, a*, r, M, p):
        return a*
    else:
        DECIDE (CALL (UPDATE-ALTERNATIVES, a, M),
                CALL (UPDATE-CUES, c, M),
                CALL (UPDATE-PARAMS, p, M),
                *RECONFIGURE (AGGREGATE-AND-ORDER,
                             ACCEPTABLE,
                             UPDATE-ALTERNATIVES,
                             UPDATE-CUES,
                             UPDATE-PARAMETERS,
                             a*, r, M, p))

```

Fig. 12: Adaptive model of heuristic decision-making. CALL is a second-order function that applies the function given as its first argument to the remaining arguments; * denotes an unpacking operator like the one in Python, i.e. the list of functions returned by RECONFIGURE is unpacked into single arguments to match the signature of DECIDE.

cluding parameters). For example an absolute threshold in ACCEPTABLE only makes sense if AGGREGATE-AND-ORDER works with cue values, not rankings. The retrieval of alternatives and cues influence each other, depending on the underlying memory structure. Mov-




Fig. 13: Interdependence of function parameters.

ing the attention to a new cue may also lead to new ideas for further alternatives and vice versa.

Many controversies in current research are encapsulated in this generalization. A configuration of UPDATE-functions, AGGREGATE-AND-ORDER and ACCEPTABLE corresponds to one heuristic, therefore the dynamic choice of alternatives corresponds to a choice of heuristics. RECONFIGURE can be used to choose among a fixed set of configurations (Marewski and Mehlhorn, 2011) or it can use the full parameter space (Glöckner et al., 2014). In our view, the choice of the parameters is simply another decision to make. RECONFIGURE is supposed to produce a vector of functions for the next iteration. It can do so by calling DECIDE, generating alternatives of vectors of the functions and evaluating them. Only at some point, the recursive call of DECIDE will have to stop, so some instance of RECONFIGURE will have to be fixed.

6 Discussion and Conclusion

We have presented a unifying computational model of decision making, covering descriptive theories of human decision making as well as computational theories of artificial systems (Simon, 1996). The framework is intended as a better means to explore, model and simulate different strategies. The instantiation options of the functions are far from trivial. In this paper we have given some ideas of how to represent cue functions and how to aggregate different cues using methods from computational social choice. Some of these methods resemble the aggregation methods in the literature on heuristic decision making.

6.1 Learning

Our model does not contain an explicit learning functionality, but offers different options of adaptation. For example, in iterative procedures with one cue per iteration, the order of the cues can be learned on the

outcome of previous decisions. Simple decision methods, such as aggregation by voting, may also serve as a bridge to more refined methods, such as weighted sums: for a new task the cruder method can be applied, while each decision provides more experience, from which, in the long run, weights or other parameters for more sophisticated methods can emerge.

Svenson (1979, p. 92) argues in the opposite direction. He proposes that over time, decision makers develop simplifying methods and apply the more sophisticated methods in new or important situations. The two views can be reconciled in that decision makers might learn over time how to configure the decision process and to know in which situations more costly decision methods pay off (Rieskamp and Otto, 2006). Svenson (1979, p. 93) notes that “[...] a decision problem may be more fruitfully viewed as a problem of classification”. This idea corresponds to a choice by recognition (Klein, 2017), which is in compassed in our model, as discussed in Section 2.3.

The model currently does not explicitly include feedback from observations, which would be necessary for reinforcement learning. We have implicitly assumed that the state of the environment and of memory is globally accessible to all the functions. With this assumption also reinforcement learning would be possible. However, the assignment of environment observations to model configurations is usually difficult. For example, in the Traveling Salesperson Problem we would first need a representation of situations rather than problem instances. Having made a single choice gives only limited feedback (if we observe that the last edge we added has led to a crossing, we can infer a negative reward, but in most situations we have no such feedback). Reward assignment is a well-known challenge of reinforcement learning and would have to be tackled in the decision making context (Rieskamp and Otto, 2006).

6.2 Relation to Cognitive Architectures

This model makes the basic assumption that intelligent decisions and behavior, both in natural and in artificial systems, emerge from an iterative process of interacting entities (Gharajedaghi, 2011). It can be regarded as an orthogonal view to cognitive architectures: Cognitive architectures make assumptions about specific components and how they interact, where usually the components are the focus of attention. In our model, the components are as open as possible (we make no assumption of how alternatives are represented, memory is organized, or how the strategies are chosen), while we focus on the wiring as an iterative procedure.

6.3 Engineering Perspective

From an engineering point of view, the model provides a structure that may help to make decisions that are better understandable for people. A program might explicitly display the alternatives and cues considered. This can also open possibilities of mixed-initiative problem solving of collaborating humans and computers. However, parameters such as aggregation function or weights, are hard to determine for users. The different choices for the components of the model constitute new parameters that again have to be optimized for or that have to be set by people. Learning or classification approaches are necessary to relieve users and developers from this choice.

We have presented a specific example of a computational decision-making process for TSP solving. The results are not intended to be generalized to other tasks and we have by far not explored the whole spectrum of instantiations of the algorithm. One especially interesting path is to examine noncommensurable aggregation methods for cues. Compromises as in weighted sums (which are the predominant method of attribute integration in artificial intelligence) can lead to acceptable average results, but may fail in special cases. As an example, we have used the Heuristic Problem Solver for robot navigation (Kirsch, 2017), where the alternatives are possible control commands and cues contain considerations such as moving towards the goal, moving forwards (instead of side- or backwards), and staying away from obstacles. If the robot starts in a narrow space with its back to the goal position, it sometimes doesn't move at all, because the urge of moving towards the goal is canceled out by the urge to move forward and to stay away from obstacles. A noncommensurable strategy might help to make sure that the robot moves at all.

6.4 Open Questions

While our model provides a unifying view on the process of decision making, it says nothing about some essential aspects of the process: representation of situations and memory organization and retrieval, and the role of habits (Wood and Neal, 2007).

These interlinked aspects are still poorly understood. The research community should invest more effort in understanding these basic mechanisms of cognition, not just in the context of decision making. Marewski and Mehlhorn (2011) illustrate the benefits of explicitly modeling memory access to better understand human decision making.

The presented use case of the Traveling Salesperson problem is suffering from this white spot in our understanding. From our data (Rach and Kirsch, 2016) we see that people use different strategies throughout one TPS instance and change their strategies when repeating instances. But we know neither how people represent a situation in the solution process nor how they manage to retrieve only such options from memory that are most promising.

The generation of promising alternatives and cues is core to make a decision process work. Maybe because our understanding of this process is so limited, engineers are still adhering to rational methods of decision making that simply assume that all alternatives are accessible at the same time. But both our understanding and use of decision processes would be greatly improved by understanding how people manage to consider very good alternatives first (March, 1988).

In sum, our model provides a general framework for research on human heuristics as well as decision making in artificial intelligence. We have given some ideas for instantiating the framework and using it in a computational context. There are many more ways in which this model can be instantiated and used in the future.

Acknowledgements With the support of the Bavarian Academy of Sciences and Humanities.

Compliance with Ethical Standards

The author has no conflicts of interest. The research involved no human participants or animals.

References

- Ariely, D. (2010). *Predictably irrational: The hidden forces that shape our decisions*. Harper Perennial, New York.
- Banzhaf, W., Nordin, P., Keller, R., and Francone, F. (1998). *Genetic Programming — An Introduction*. Morgan Kaufmann, San Francisco, CA.
- Brandt, F., Conitzer, V., Endriss, U., Lang, J., and Proccaccia, A. D., editors (2016). *Handbook of Computational Social Choice*. Cambridge University Press.
- Dantzig, G. B. (1963). *Linear programming and extensions*. Princeton University Press.
- Gharajedaghi, J. (2011). *Systems Thinking: Managing Chaos and Complexity*. Morgan Kaufmann, 3rd edition.
- Gigerenzer, G. (2001). The adaptive toolbox. In Gigerenzer, G. and Selten, R., editors, *Bounded ratio-*

- nality: The adaptive toolbox*. MIT Press, Cambridge, MA.
- Gigerenzer, G. and Brighton, H. (2009). Homo heuristics: Why biased minds make better inferences. *Topics in Cognitive Science*, 1:107–143.
- Gigerenzer, G. and Goldstein, D. G. (1999). Betting on one good reason: The take the best heuristic. In Gigerenzer, G., Todd, P. M., and the ABC Research Group, editors, *Simple Heuristics That Make Us Smart*. Oxford University Press.
- Glöckner, A., Hilbig, B. E., and Jekel, M. (2014). What is adaptive about adaptive decision making? a parallel constraint satisfaction account. *Cognition*, 133:641–666.
- Graham, S. M., Joshi, A., and Pizlo, Z. (2000). The traveling salesman problem: A hierarchical model. *Memory & Cognition*, 28(7):1191–1204.
- Hertwig, R., Hoffrage, U., and Martignon, L. (1999). Quick estimation: Letting the environment do the work. In Gigerenzer, G., Todd, P. M., and the ABC Research Group, editors, *Simple Heuristics That Make Us Smart*. Oxford University Press.
- Huber, O. (1979). Nontransitive multidimensional preferences: Theoretical analysis of a model. *Theory and Decision*, 10:147–165.
- Katsikopoulos, K. V. and Martignon, L. (2006). Naïve heuristics for paired comparisons: Some results on their relative accuracy. *Journal of Mathematical Psychology*, 50(5):488–494.
- Keller, N. and Katsikopoulos, K. V. (2016). On the role of psychological heuristics in operational research; and a demonstration in military stability operations. *European Journal of Operational Research*, 3:1063–1073.
- Kirsch, A. (2011). Humanlike problem solving in the context of the traveling salesperson problem. In *AAAI Fall Symposium on Advances in Cognitive Systems*.
- Kirsch, A. (2012). Hierarchical knowledge for heuristic problem solving — a case study on the traveling salesperson problem. In *First Annual Conference on Advances in Cognitive Systems*.
- Kirsch, A. (2016). Heuristic decision-making for human-aware navigation in domestic environments. In *2nd Global Conference on Artificial Intelligence (GCAI)*.
- Kirsch, A. (2017). A modular approach of decision-making in the context of robot navigation in domestic environments. In *3rd Global Conference on Artificial Intelligence (GCAI)*.
- Klein, G. (2017). *Sources of Power: How People Make Decisions*. The MIT Press. MIT Press.
- Kotseruba, I. and Tsotsos, J. K. (2016). A review of 40 years of cognitive architecture research: Core cognitive abilities and practical applications. arXiv:1610.08602.
- Langley, P. (2017). Progress and challenges in research on cognitive architectures. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*.
- Lee, M. D. and Newell, B. R. (2011). Using hierarchical bayesian methods to examine the tools of decision-making. *Judgment and Decision Making*, 6(8):832–842.
- MacGregor, J. and Ormerod, T. (1996). Human performance on the traveling salesman problem. *Perception & Psychophysics*, 58(4):527–539.
- MacGregor, J. N. and Chu, Y. (2011). Human performance on the traveling salesman and related problems: A review. *The Journal of Problem Solving*, 3(2).
- March, J. G. (1988). *Bounded rationality, ambiguity, and the engineering of choice*, pages 33–57. Cambridge University Press.
- Marewski, J. N. and Mehlhorn, K. (2011). Using the act-r architecture to specify 39 quantitative process models of decision making. *Judgment and Decision Making*, 6(6):439–519.
- Marewski, J. N. and Schooler, L. J. (2011). Cognitive niches: An ecological model of strategy selection. *Psychological Review*, 118(3):393–437.
- Newell, A. (1982). The knowledge level. *Artificial Intelligence*, 18:81–132.
- Newell, A. and Simon, H. (1972). *Human Problem Solving*. Prentice Hall, Upper Saddle River, New Jersey.
- Rach, T. and Kirsch, A. (2016). Modelling human problem solving with data from an online game. *Cognitive Processing*, 17(4):415–428.
- Rieskamp, J. and Hoffrage, U. (1999). When do people use simple heuristics and how can we tell? In Gigerenzer, G., Todd, P. M., and the ABC Research Group, editors, *Simple Heuristics That Make Us Smart*. Oxford University Press.
- Rieskamp, J. and Otto, P. E. (2006). Ssl: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, 135(2):207–236.
- Russell, S. and Norvig, P. (2003). *Artificial Intelligence - A Modern Approach*. Prentice Hall, Upper Saddle River, New Jersey.
- Russo, J. E. and Doshier, B. A. (1983). Strategies for multiattribute binary choice. *Journal of Experimental Psychology: Learning, Memory, Cognition*, 9(4):676–696.
- Shah, A. K. and Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, 134(2):207–222.

- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2):129–138.
- Simon, H. A. (1996). *The Sciences of the Artificial*. MIT Press, 3rd edition.
- Svenson, O. (1979). Process descriptions of decision making. *Organizational Behavior and Human Performance*, 23:86–112.
- Tak, S., Plaisier, M., and van Rooij, I. (2008). Some tours are more equal than others: The convex-hull model revisited with lessons for testing models of the traveling salesperson problem. *The Journal of Problem Solving*, 2(1).
- Tenbrink, T. and Wiener, J. (2009). The verbalization of multiple strategies in a variant of the traveling salesperson problem. *Cognitive Processing*, 10:143–161.
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131.
- Wolter, D. and Kirsch, A. (2015). Leveraging qualitative reasoning to learning manipulation tasks. *Robotics*, 4(3):253–283.
- Wood, W. and Neal, D. (2007). A new look at habits and the habit-goal interface. *Psychological Review*, 114(4):843.
- Zwicker, W. S. (2016). *Introduction to the Theory of Voting*, chapter 2. Cambridge University Press.