

A glass-box interactive machine learning approach for solving NP-hard problems with the human-in-the-loop

ANDREAS HOLZINGER^{1,2}, MARKUS PLASS^{1,2}, KATHARINA HOLZINGER^{1,2},
GLORIA CERASELA CRIȘAN³, CAMELIA-M. PINTEA⁴ and VASILE PALADE⁵

ABSTRACT. The ultimate goal of the Machine Learning (ML) community is to develop algorithms that can automatically learn from data, to extract knowledge and to make decisions without any human intervention. Specifically, automatic Machine Learning (aML) approaches show impressive success, e.g. in speech/image recognition or autonomous drive and smart car industry. Recent results even demonstrate intriguingly that deep learning applied for automatic classification of skin lesions is on par with the performance of dermatologists, yet outperforms the average human efficiency. As human perception is inherently limited to 3D environments, such approaches can discover patterns, e.g. that two objects are similar, in arbitrarily high-dimensional spaces what no human is able to do. Humans can deal simultaneously only with limited amounts of data, whilst "big data" is not only beneficial but necessary for aML. However, in health informatics, there are few data sets; aML approaches often suffer from insufficient training samples. Many problems are computationally hard, e.g. subspace clustering, k-anonymization, or protein folding. Here, interactive machine learning (iML) could be successfully used, as a human-in-the-loop contributes to reduce a huge search space through heuristic selection of suitable samples. This can reduce the complexity of NP-hard problems through the knowledge brought in by a human agent involved into the learning algorithm. A huge motivation for iML is that standard black-box approaches lack transparency, hence do not foster trust and acceptance of ML among end-users. Most of all, rising legal and privacy aspects, e.g. the European General Data Protection Regulations (GDPR) make black-box approaches difficult to use, because they often are not able to explain why a decision has been made, e.g. why two objects are similar. All these reasons motivate the idea to open the black-box to a glass-box. In this paper, we present some experiments to demonstrate the effectiveness of the iML human-in-the-loop model, in particular when using a glass-box instead of a black-box model and thus enabling a human directly to interact with a learning algorithm. We selected the Ant Colony System (ACS) algorithm, and applied it on the Traveling Salesman Problem (TSP). The TSP-problem is a good example, because it is of high relevance for health informatics as for example on protein folding problem, thus of enormous importance for fostering cancer research. Finally, from studies of learning from observation, i.e. of how humans extract so much from so little data, fundamental ML-research also may benefit.

1. INTRODUCTION AND MOTIVATION FOR RESEARCH

Automatic Machine Learning. The ultimate goal of the machine learning (ML) community is to develop algorithms/systems which can automatically learn from data, extract knowledge and make predictions and decisions without any human intervention [56]. Automatic Machine Learning (aML) approaches have made enormous advance and have had practical success in many different application domains, e.g. in speech recognition [23], recommender systems [11], or autonomous vehicles [10], and many other industrial applications.

Recently, deep learning algorithms, supported by supercomputers, cloud-CPU's and extremely large data sets [17] have demonstrated to exceed human performance in visual tasks, particularly on playing games such as Atari 2600 collection [43], or playing Go [58].

Received: 22.01.2019. In revised form: 14.05.2019. Accepted: 21.05.2019

2010 *Mathematics Subject Classification.* 90C35, 05C85, 68T20, 68T05.

Key words and phrases. *Machine Learning, Traveling Salesman problem, NP-hard, Ant Colony Optimizaton.*

Corresponding author: Camelia-M. Pintea; dr.camelia.pintea@ieee.org

An impressive example from the medical domain is the recent work by Esteva et al. [20]: they utilized a GoogleNet Inception v3 CNN architecture [62] for the classification of skin lesions using a single CNN, trained end-to-end from images directly, using only pixels and disease labels as inputs. They pre-trained their network with 1.28 million images (1,000 object categories) from the 2014 ImageNet Challenge [54], and trained it on 129,450 clinical images, consisting of 2,032 different diseases. The performance was tested against 21 board-certified dermatologists on biopsy-proven clinical images with two critical binary classification use cases: keratinocyte carcinomas versus benign seborrheic keratoses; and malignant melanomas versus benign nevi. The results show that CNN achieves a performance on par with all tested human experts across both tasks, demonstrating the ability of classifying skin cancer by ML with a level of competence comparable to dermatologists. This work proved that automatic machine learning (aML) works well when large amounts of training data are available [59]; therefore, big data management, which is often considered as a burden, is here not only beneficial but necessary.

Disadvantages of black-box approaches. Besides of being resource intensive and data hungry, black-box approaches have, at least when applied to safety-critical domains such as the medical domain, one enormous drawback: by lacking transparency, they often are not able to explain why a decision has been made. This does not foster trust and acceptance among end-users. Most of all, legal aspects make black-box approaches difficult: after the European General Data Protection Regulations (GDPR) took effect on June, 1st, 2018, customers are given a right-to-be-forgotten [40], i.e. having their data deleted on request. Whilst ensuring privacy is both positive and necessary, this could lead to a competitive disadvantage for data-driven and data-dependent European companies. Consequently, it becomes important to understand what effects the removal of certain data will have on the performance of ML techniques. Moreover, the need for privacy-aware ML pipelines and the production of k-anonymized open data sets for real-world usage will become urgent [41].

Representation Learning and Context. The performance of ML algorithms is dependent on the choice of the data representations. Current ML algorithms are still unable to extract the discriminative knowledge from data. Bengio et al. [6] argue that this can only be achieved if the algorithms can learn to identify and to disentangle the underlying exploratory factors already existent among the low-level data. This supposes that a truly intelligent algorithm is required to understand the context, and to be able to discriminate between relevant and irrelevant features similarly as we humans can do. The questions What is interesting? and What is relevant? are inherently hard questions, and as long as we cannot achieve true intelligence with automatic approaches, we have to develop algorithms which can be applied by a human expert. Such a domain expert is likely to be aware of the context. Following the probabilistic perspective, this would mean that learning features from data can be seen as recovering a parsimonious set of latent random variables (i.e., according to Occam's razor, see [18] for a critical discussion), representing a distribution over the observed data to express a probabilistic model $p(x, h)$ over the joint space of the latent variables, h , and the observed data x . This approach fits well into the perspective of Cognitive Sciences [68].

Motivation for a human-in-the-loop. Interactive Machine Learning (iML) has various definitions and is often referred when users face machine learning approach [4]. Other definitions speak also of a human-in-the-loop, but it is what we would rather call classic supervised ML approaches [57], and a total different meaning is to include humans into a feedback loop [55].

By integrating a human-in-the-loop (e.g., a human kernel [68]), or the involvement of a human directly into the algorithm, iML is defined by algorithms which interact with agents and can optimize their learning behaviour through this interaction where the agents can be humans [26]. The general idea is making use of the strengths of human cognitive abilities -when automatic approaches fail. Consequently, iML-approaches can be of particular interest to solve problems, where we are lacking big data sets, dealing with complex data and/or rare events, where aML suffer of insufficient training samples.

In the medical domain a doctor-in-the-loop can help with his/her expertise in solving problems which otherwise would remain NP-hard. A recent experimental work [26] demonstrates the usefulness on the Traveling Salesman Problem (TSP), which appears in a number of practical problems, e.g., the native folded three-dimensional conformation of a protein in its lowest free energy state; both 2D and 3D folding processes are assumed to be conditionally intractable [13]. TSP is about finding the shortest circuit through a set of points; it is an intransigent mathematical problem, having many heuristic solving methods [38]. There is evidence that the inclusion of a human can be useful in numerous other problems in different application domains, see [3, 45]. However, for clarification, iML means the integration of a human into the algorithmic loop, i.e., to open the black box approach to a glass box.

The automatic machine learning (aML) has not always good results when solving complex problem; Interactive Machine Learning (iML) could improve the efficiency of solving NP-hard problems by involving the human knowledge.

The algorithms within iML include interactions with different agents including human ones [26]. Figure 1 [29, 30] illustrates the iML model with the role of the human in the algorithms loop and show the difference from the classic supervised learning.




FIGURE 1. In the iML human-in-the-loop model the humans are involved in many tasks as: pre-processing and selecting data/features and directly into the learning phase, interacting with the algorithm; this last task transforms into a glass-box the initial black-box approach. The iML approach allows crowd-sourcing or other related approaches [29, 30].

In medicine, in particular in diagnosing radiological images, a high-level expert (expert in-the-loop) [2] use his knowledge on interpretation of complex patterns into the retrieval process. In some cases doctors diagnose without explicit rules, based on their intuition; iML could enhance algorithms with knowledge-based intuition. In many domains, including medicine, to find a feasible solution in short time is a must; in a clinical examination, a doctor has just few minutes to make a decision [21], to solve a difficult problem. So, an approximate solution to a difficult problem in a short time, as an heuristic provides, is rather beneficent.

Nowadays, meta-heuristics are known to successfully solve combinatorial optimization problems. Ant Colony Optimization (ACO) is a bio-inspired meta-heuristic which offer several ant-based algorithms used to solve different real-life problems as for example the Linear Ordering Problem [12], the Matrix Bandwidth Minimization Problem [50],

the gate assignment problem [48], the fixed-charged transportation problem [51], routing problems including versions of Traveling Salesman Problem (TSP) as the railway TSP [49].

Learning from few examples. A further motivation for our research is the fact that humans successfully learn from few examples. A good example is Gaussian process, where aML approaches (e.g., kernel machines [24]) struggle on function extrapolation problems, which are quite trivial for human learners [22]. As one of the NAE Grand 21st Century challenges is health informatics, this paper is a step towards the knowledge enhancement of this effervescent research domain [1].

Future information systems will have to address the balance between personal data privacy and seamless data sharing. Ubiquitous computing will improve health care and will pose new challenges to the networking systems. Future robust computerized systems will survey and react to public health emergencies both at local and global scale.

The current article discusses if a direct intervention of a human agent in an algorithm improves algorithm's results. Discussions are based on tests and comparisons results from the iML-project's web-platform [32]. The human-into-the-loop could solve difficult problems from different domains including medicine (doctor-into-the-loop [33, 34]) by using for example collaborative interactive ML [52].

2. BACKGROUND AND RELATED WORK

2.1. Problem Solving: Human versus Computer. Nowadays ML has not yet too much success on performing extrapolation problems as humans do. When asked to make extrapolations of functions drawn from Gaussian Processes (GP) with known kernels in sequence, humans progressively learn about the kernel [68].

We do not know yet how do humans face new situations [70]. Humans use inductive reasoning and generalize from few examples using Bayesian inference [64]. Let $X = \{x_1, \dots, x_n\}$ be a set of n observations from a novel entity C (e.g. a word) and the examples given are from known entities U . It is assumed that the human learner has access to a hypothesis space $H = \{H_1, H_2, \dots, H_n\}$, and a probabilistic model relating each hypothesis $h \in H$ to X . Simplified, each h can be seen as a pointer to a subset of entities in the domain that is a candidate extension for C . The learner is assumed to be able to identify the extension of each h . Given the examples X , the Bayesian learner evaluates all hypotheses for candidate meanings according to Bayes rule, by computing their posterior probabilities $p(h|X)$, proportional to the product of prior probabilities $p(h)$ and likelihoods $p(X|h)$:

$$p(h|X) = \frac{p(X|h)p(h)}{\sum_{h' \in H} p(X|h')p(h')} \propto p(X|h)p(h) \quad (2.1)$$

The prior $p(h)$ represents the learners expectation about plausible meanings for entity C and is independent of the observed examples X and reflects conceptual constraints, different contexts, or beliefs conditional on the meanings of other previously learned entities.

The likelihood $p(X|h)$ captures the statistical information inherent in the examples X , and reflects expectations about which entities are likely to be observed as examples of C given a particular hypothesis h about the meaning of C .

The posterior $p(h|X)$ reflects the learners degree of belief that h is in fact the true meaning of C , given a combination of the observations X with prior knowledge about plausible word meanings.

Whilst such a Bayesian approach provides a powerful computational framework for explaining how humans solve inductive problems, much open questions are remaining [65, 9, 60].

Why is this important for machine learning? A true computationally intelligent algorithm will take its own decisions after extracting knowledge from automatically learning. Consequently, machine learning efforts have always been inspired by how humans learn, extract knowledge and make decisions [69]. From a probabilistic perspective, the ability for a model to automatically discover patterns and perform extrapolations is determined by an existing solution. Such a model should represent many possible solutions to a given problem with inductive biases, which can extract intricate structures from limited data. In this context function learning is central for cognition: the tasks require the construction of mental representations that map inputs to outputs $f : X \rightarrow Y$. Humans estimate functions or perform associative learning. Lucas et al. [37] developed a rational model of human function learning that combines the strengths of both these theories.

Human Abilities on Optimization Problems. Getting insight into human abilities involved in optimization problems is considered of raising interest in the next decade [39] as it could increase our knowledge about visual representation and its support. The involved, relatively low complexity could be utilized to help to solve several optimization problems. Most of all, it could help to solve hard problems, because many computational tasks of practical interest and of relevance to the health informatics domain are exceptionally difficult to solve [44].

Macgregor and Ormerod ([38]) compared human performance on the Traveling Salesman Problem (TSP) to heuristic algorithms. They concluded that heuristics provided little to explain human performance, and that human problem solvers instead used a perceptual process to solve the problem; this was confirmed by two experiments carried out by Best and Simon [8]. They collected sequences of mouse moves and clicks made by humans when solving TSP instances. Their first experiment replicated the problem set used by Macgregor and Ormerod [38]. The second experiment used problems with uniform random distributions of points. Mouse movements were examined at the level of selection of individual moves (clicks). The TSP instances were presented one at a time on a computer screen. Each problem started with a dialog box centered on the screen asking whether the participant was ready to start; by clicking the OK button, the mouse was also adjusted on exactly the center of the display. The participants then used the mouse to select nodes until completing a tour. They could not backtrack or undo moves. The results indicated that the Nearest Neighbor algorithm best fit the human TSP solution methods. Although the overall solution obtained using a Nearest Neighbor algorithm did not produce solutions of the same quality as human problem solvers, a majority of the moves made by human problem solvers are to the closest point. When the local processing constraints suggested by the Nearest Neighbor algorithm are combined with a global plan for the general shape of the solution, a simulation of humans was quite effective.

There is evidence that humans are very good in finding feasible solutions to many decision problems under uncertainty, and specifically to the TSP. Moreover, humans do not realize the computational complexity of the problem, but are able to exploit structural properties of the overall instance to improve parts of it [35, 63].

Acuña and Parada [1] tested the solutions of $N = 28$ participants to $M = 28$ instances of the Euclidean Traveling Salesman Problem [53]. In their experiment the participants were provided feedback on the cost of their solutions and were allowed unlimited solution attempts (as many trials as they liked). The authors found a significant improvement between first and last trials, the participants' modified their current solutions as the good edges were identified showing a good structural exploitation. More trials harmed the participants' ability to separate good from bad edges, suggesting that after too many trials the participants simply ran out of ideas.

Human vs. Computer. In image optimization and other large data automatic processes of the aML algorithms frequently outperform humans, but in other processes as for example in natural language translation or other complex problems aML algorithms fail. These facts are based on the intrinsic characteristics of a machine: a machine it is not “intuitive” and naturally not “creative”. Recently, the world Go champion Lee Sedol won a match from five when confronting AlphaGo (Google DeepMind team). So, the human’s creativity could beat machines procedures and rules [25].

3. ANT COLONY SYSTEM SOLVING THE TRAVELING SALESMAN PROBLEM

Traveling Salesman Problem (TSP) is a classical NP-hard problem in computer science with relevance also for the health domain. TSP is used to test the introduced iML model. The TSP has been studied using different techniques including meta-heuristics as the Ant Colony Optimization (ACO) to find beneficent solutions [14, 15].

The Ant Colony System is an Ant Colony Optimization [19] version that starts from an initial solution to the TSP instance and iteratively improves it. The initial solution can be randomly generated or other rapid construction method can be used. A constant number of m ants form the ant colony that evolves and searches the shortest Hamiltonian tour in a complete weighted graph with n vertices. The pheromone on edges form the matrix τ_{ij} .

The initial pheromone quantity on each edge is constant. Three main activities form an iteration can overlap. The solution construction is based on the pseudo-random-proportional rule, which is a combination of deterministic selection of the most promising vertex, and a probabilistic process of choosing the next vertex. The parameters that controls the solution constructions are: β , the weight of the heuristic function η on edges (which is analytically expressed by the inverse of the distances) and q_0 , the threshold for deterministic decisions.

If a specific ant stays in vertex i , the available vertices form the set J_i . The agent randomly generates a value $q \in [0, 1]$. If $q \leq q_0$, then the next vertex j is deterministically set using formula 3.2.

$$j = \underset{l \in J_i}{\operatorname{argmax}} (\tau_{il} \cdot [\eta_{il}]^\beta) \quad (3.2)$$

Otherwise, j is chosen using the transition probabilities from formula 3.3. After all the solutions are constructed, their lengths are computed and if a better solution was found, then the overall best solution P_{best} is updated.

$$p_{ij} = \frac{\tau_{ij} \cdot [\eta_{ij}]^\beta}{\sum_{l \in J_i} \tau_{il} \cdot [\eta_{il}]^\beta}, j \in J_i \quad (3.3)$$

Each ant updates the pheromone on its path using equation 3.4 where $\rho \in [0, 1]$ is the evaporation parameter and $L_{initial}$ is the length of the initial tour.

$$\tau_{ij}(t+1) = (1 - \rho) \cdot \tau_{ij}(t) + \rho \cdot \frac{1}{n \cdot L_{initial}} \quad (3.4)$$

The current iteration ends by reinforcing the pheromone on the P_{best} edges using the formula 3.5, where L_{best} is the length of P_{best} . The iterations are repeated until a stopping criterion is met.

$$\tau_{ij}(t+1) = (1 - \rho) \cdot \tau_{ij}(t) + \rho \cdot \frac{1}{L_{best}} \quad (3.5)$$

Is a current practice to use local improvement methods [16, 7, 36] after each ant completes its tour. Other ant-based improvements were on transition policy [47]. ACS was further developed. For example, the Inner Ant System introduces a pheromone update rule for improving the ants’ local search [46].

Although initially designed as multi-agent systems that evolve independently, without central guidance, new ACO algorithms focus on exploiting more the information linked to good solutions or promising regions from the search space and avoiding to be kept in a local optimum solution. This guidance is designed as an activity that monitors the colony activity and reacts when it is needed. One such example is Foot Stepping technique described in [67]. It is meant to decrease the pheromones on the best paths when stagnation manifests and thus to increase the exploration. Adaptation is managed also by including pheromone alteration as in [5]. When a dynamic TSP with sudden modifications of the costs is approached, the pheromone values are modified. It is assumed that the algorithm knows when changes appear and immediately reacts.

All the guided actions are modules that automatically react when some pre-set criterion is met. To our knowledge, this paper is the first attempt to introduce totally external guidance to a working, solving process. Our investigation was founded on the good ability of humans to solve TSP instances. There are many studies in the literature that describe the process of trained or untrained TSP solving. For a comprehensive review see [39].

The humans develop the ability to better solve an instance on repeated trials, they perform better on visually described instances with fewer vertices, they rarely provide crossing tours, and instinctively use the convex hull property (the vertices on the convex hull are traversed in their order). A visual interface for a game designed for assessing the human abilities for solving small TSP 2D instance is presented in [1].

4. INTRODUCING NEW CONCEPTS BASED ON HUMAN INTERACTION

As described in Holzinger et al. [28], the human interaction is based on changing the ants transition rules. As is proposed as a *Future outlook* section of our previous work [30], here we re-iterate two novel concepts: Human-Interaction-Matrix (HIM) and Human-Impact-Factor (HIF).

- **Human-Interaction-Matrix (HIM).** The creation of a HIM allows the human to control the ants actions. The user sets the probabilities for performing her/his decisions. These values are dynamically interpreted by each ant, during its solution construction. The matrix is dynamically modified by the users decisions.
- **Human-Impact-Factor (HIF).** The HIF is the variable interpretation of the HIM by each ant. An example is presented in Table 1. When an ant is in node B, there are half chances to be sent to C, based on HIM, else it moves using transition rules and C is removed from the destinations' set.

When an ant is in node C, there are half chances to be sent in B and 10% chances to be sent to E. If the ant does not go to B or E, then the implemented algorithm decides where to go, excluding B and E from the available nodes. All these interactions have no direct influence on the pheromones on the edges; they are just external guidance for the ants. If the user decides that one edge is no longer of interest, the corresponding HIM can be set to zero and the ants will not be influenced anymore.

4.1. Proposed iML Ant Colony Algorithm. In the new algorithm ACS-iML, the *ConstructionSolution* method works has also the parameters HIM and HIF:

```
if ( $q \leq HIF$ ) then HIM identifies next move
else ants identify next move
```

where q is an randomly generated number between $[0, 1]$.

In Figure 3 a draft of the GUI is presented. The upper part of the machine-learning section contains all information from the standard Ant-Algorithm in pseudocode (Figure 2).

Algorithm: Ant Colony Algorithm iML

Input : ProblemSize, m , β , ρ , σ , q_0

Output: P_{best}

$P_{best} \leftarrow \text{CreateHeuristicSolution}(\text{ProblemSize});$
 $P_{best_cost} \leftarrow \text{Cost}(P_{best});$
 $\text{Pheromone}_{init} \leftarrow \frac{1.0}{\text{ProblemSize} \times P_{best_cost}};$
 $\text{Pheromone} \leftarrow \text{InitializePheromone}(\text{Pheromone}_{init});$

while $\neg \text{StopCondition}()$ **do**

for $i = 1$ **to** m **do**

$S_i \leftarrow \text{ConstructSolution}(\text{HIM}, \text{HIF}, \text{Pheromone}, \text{ProblemSize}, \beta, q_0);$
 $S_{i_cost} \leftarrow \text{Cost}(S_i);$
 if $S_{i_cost} \leq P_{best_cost}$ **then**

$P_{best_cost} \leftarrow S_{i_cost};$
 $P_{best} \leftarrow S_i;$

end
 $\text{LocalUpdateAndDecayPheromone}(\text{Pheromone}, S_i, S_{i_cost}, \rho);$

end
 $\text{GlobalUpdateAndDecayPheromone}(\text{Pheromone}, P_{best}, P_{best_cost}, \rho);$
 while $\text{isUserInteraction}()$ **do**

$\text{UpdateHumanInteractionMatrix}(\text{HIM});$

end

end
return $P_{best};$

FIGURE 2. Pseudocode of the Ant Colony System iML algorithm

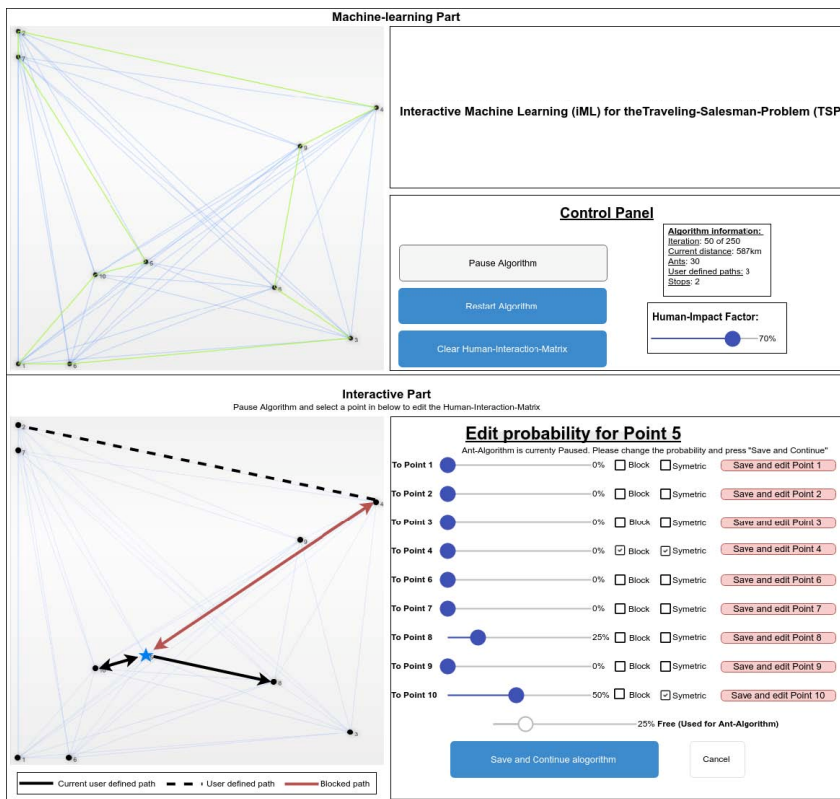


FIGURE 3. Draft of the GUI

The current shortest circuit is marked in green. All edges have the same width, a distinct feature from the version described in [28]. The HIF can be adjusted by a single slider, during the whole run time of the algorithm.

The bottom part of the GUI allows the user to change the Human-Impact-Matrix. By clicking on one of the nodes the edit-mask of the Points opens. The mask now allows adjusting the probability for taking the corresponding edge (as described above). For each Point 100% can be spread over the other points.

An edge could also be locked completely by clicking the Block Button. If this is done, the rest of the pheromones will not be spread towards this Point. All changes made by humans are summarized in an internal picture (see Figure 3, bottom left). So the workflow can be summarized as:

- Selection of the HIF
- Start and initialization of the Ant-Algorithm
- Pause the Algorithm
- Selection of Node
- Modification of the HIM
- Save the changed HIM
- Another click on Pause/Resume continues the algorithm.
- These steps are repeated as many times as needed.

The implementation is browser-based (Java-Script); it could be used on any operating system and has all web-based applications advantages. The parameter values were: 30 ants, 250 iterations, $\beta = 3$, $\rho = 0.1$.

Figure 3 illustrates the current tour with a green line and the best tour with a red line. Pre-defined data sets [66] are used to easily test the application; in the example the burma14.tsp data-set is included. Comparison of the solutions with the optimal ones is provided by clicking on the Compare with optimal tour button.

5. EXPERIMENTAL METHOD, SETTING AND RESULTS

As an aML algorithm, ACO does not usually interact with the computing environment. The ants walk around and update the artificial pheromone on edges after iteration. This procedure is repeated until a stopping criterion is met. Following the iML-approach, the human now can open the ACO black-box and can manipulate this algorithm by dynamically changing the behavior of the ants. This is done by modifying the transition rules (3.2, 3.3) from the Ant Colony System structure.

	A	B	C	D	E		A	B	C	D	E
A	0	0	0	0	0	A	N	0	0	0	0
B	0	0	0	0	0	B	0	N	0.5	0	0
C	0	0	0	0	0	C	0	0.5	N	0	0
D	0	0	0	0	0	D	0	0	0	N	0
E	0	0	0	0	0	E	0	0	0.1	0	N

a)
b)

TABLE 1. Example of Human-Interaction-Matrix (HIM) and Human-Impact-Factor (HIF). Initially, no human intervention. The Human user sets 50% chances to go from B to C and vice versa, and 10% chances to go from C to E. In the other cases, the solver proceeds on its own. The loops are forbidden (N on the main diagonal).




FIGURE 4. Example of Human interaction in guiding ants, based on Table 1. Initially the matrix is empty; the Human-Impact-Factor (HIF) value $HIF=0.5$ (blue squares) force the ants to go from B to C in 50% of the cases, and from C to B in 50% of the cases. The Human-Impact-Factor (HIF) value $HIF=0.1$ (green square) forces the ants to go from C to E in 10% of the cases. Note the symmetric set for the pair (B C) and the asymmetric set for (C E). In all the other cases, the ants move based on the transition rule only to the nodes having zero values (light blue squares). The loops are excluded.

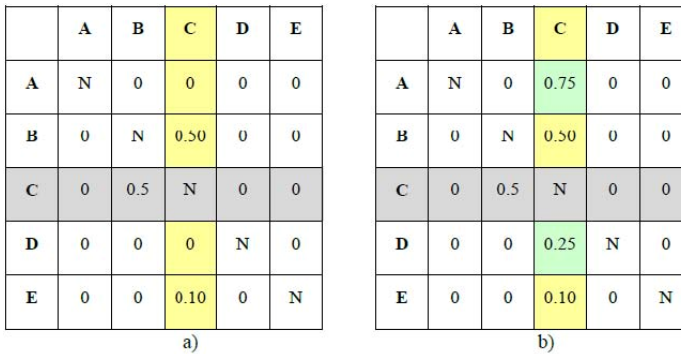


FIGURE 5. a) Step 1: the ant is in C, line C is locked (Grey background). b) Step 2, case 3: the ant ignores the human and computes the probabilities for going to A or D (green cells).

The current section illustrates several transition cases. As the current work considers asymmetric paths, the HIM matrix can also be asymmetric. It follows the description of a step of the algorithm starting from the matrix in Table 1 b) and visually illustrated in Figure 4. Let us suppose that an ant is initially dispatched in node C. In Figure 5 a), the line C is locked (Grey background), meaning that C is unavailable for further moves.

Let us suppose that the transition rules (3.2, 3.3) gave the following probabilities for moving without human intervention to available nodes: 30% to A, 40% to B, 10% to D and 20% to E. The ant uniformly generates a random value q_0 . We have the following cases, depending on which interval q_0 falls:

- Case 1 If $q_0 \in [0, 0.5]$ then the ant follows the human decision and goes to node B.
- Case 2 If $q_0 \in [0.5, 0.6]$ then the ant follows the human decision and goes to node E.
- Case 3 If $q_0 \in [0.6, 1]$ then the ant follows the transition rules (3.2, 3.3). As in this case the nodes B and E are forbidden, the available nodes are A and D (zeroes in Fig. 5 a)

and marked with green in Fig. 5 b)). As the transition without human intervention from C to A had the probability 30%, the transition from C to D had to be made with 10% probability, and only these two nodes are allowed, the ant normalizes these two values and computes new transition probabilities: 75% for going to node A and 25% for moving to D. The ant moves accordingly either to A or to D.

No matter where the ant just arrived, the corresponding line is locked and the process is repeated until the ant constructs a complete tour.

6. DISCUSSION

With the creation of Human-Impact-Factor and Human-Interaction-Matrix, the human decisions moved the algorithm control from Pheromone-Based to Ant-Based. This allows that the pheromones only are modified by the ants not by the human, so a possible flooding with pheromones can be avoided.

The current algorithm and the version from [28] could be further extended as local search methods. This is because they work better for a small number of nodes.

As an example we have included Figure 6, with two copies of *burma14.tsp* [66] instance, so an instance with two clusters of 14 nodes each. After solving each cluster, the solution for the new instance with 28 nodes is based on the partial solutions, connected in a less costly way.

A larger instance should be clustered and the local results could be used to solve the larger instance. The increase of the cities number is also possible as in the preprocessing phase of subspace clustering [31, 42, 61].

7. CONCLUSION

The paper shows how iML model [27] including human-computer interaction successfully improves the quality of the solutions to the classical Traveling Salesman Problem. Furthermore, other NP-hard problems will be in our attention to be solved using human-computer interactions [4]. We will focus not only on medical-related problems but also from other domains where the human intervention will make an important impact on problem's solutions.




FIGURE 6. The iML human-in-the-loop approach as local search for an instance divided in two clusters and each one is solved with iML algorithm. In this case two short connections between the two circuits could provide a solution to the initial instance with 28 nodes.

Acknowledgments. We thank our international colleagues from the HCI-KDD and IET-CIS-IEEE networks for valuable discussions.

REFERENCES

- [1] Acuña, D. E. and Parada, V., *People efficiently explore the solution space of the computationally intractable traveling salesman problem to find near-optimal tours*, PloS one, **5** (2010) No. 7, 1–10
- [2] Akl, C. B. et al., *Content-based image retrieval in radiology: Current status and future directions*, Journal of Digital Imaging, **24** (2011), No. 2, 208–222
- [3] Amato, R. et al., *A multi-step approach to time series analysis and gene expression clustering*, Bioinformatics, **22** (2006), No. 5, 589–596
- [4] Amershi, S., Cakmak, M., Knox, W. B. and Kulesza, T., *Power to the people: The role of humans in interactive machine learning*, AI Magazine, **35** (2014), No. 4, 105–120
- [5] Angus, D. and Hendtlass, T., *Dynamic ant colony optimisation*, Applied Intelligence, **23** (2005), No. 1, 33–38
- [6] Bengio, Y., Courville, A. and Vincent, P., *Representation learning: A review and new perspectives*, IEEE transactions on pattern analysis and machine intelligence, **35** (2013), No. 8, 1798–1828
- [7] Bentley, J. J., *Fast algorithms for geometric traveling salesman problems*. ORSA Journal on computing, **4** (1992), No. 4, 387–411
- [8] Best, B. J. and Simon, H. A., *Simulating human performance on the traveling salesman problem*, Conference on Cognitive Modeling, (2000), 42–49
- [9] Chater, N. Tenenbaum, J. B. and Yuille, A., *Probabilistic models of cognition: Conceptual foundations*, Trends in cognitive sciences, **10** (2006), No. 7, 287–291
- [10] Chen, C., Seff, A., Kornhauser, A. and Xiao, J., *Deepdriving: Learning affordance for direct perception in autonomous driving*, IEEE International Conference on Computer Vision, (2015), 2722–2730
- [11] Cheng, H-T., et al., *Wide and deep learning for recommender systems*, ACM Proceedings of the 1st Workshop on Deep Learning for Recommender Systems, (2016), 7–10
- [12] Chira, C. et al., *Solving the Linear Ordering Problem using Ant Models*, GECCO 2009, ACM (2009), 1803–1804
- [13] Crescenzi, P., Goldman, D., Papadimitriou, C., Piccolboni, A. and Yannakakis, M., *On the complexity of protein folding*, Journal of computational biology, **5** (1998) No. 3, 423–465
- [14] Crişan, G-C., Nechita, E. and Palade, V., *Ant-based system analysis on the traveling salesman problem under real-world settings*, Combinations of Intelligent Methods and Applications, Springer, (2016), 39–59
- [15] Crişan, G-C., Pinteau, C-M. and Palade, V., *Emergency management using geographic information systems: application to the first Romanian traveling salesman problem instance*, Knowledge and Information Systems, **50** (2017), No. 1, 265–285
- [16] Croes, G. A., *A method for solving traveling-salesman problems*, Operations research, **6** (1958), No. 6, 791–812
- [17] Deng, J., Dong, W., Socher, R., Li, L-J., Kai, L. and Li, F-F., *Imagenet: A large-scale hierarchical image database*, IEEE Conference on Computer Vision and Pattern Recognition, (2009), 248–255
- [18] Domingos, P., *The role of occams razor in knowledge discovery*, Data Mining and Knowledge Discovery, **3** (1999), No. 4, 409–425
- [19] Dorigo, M. and Stützle, T., *Ant Colony Optimization*, MIT Press. First Edition, 2004
- [20] Esteva, A., et al., *Dermatologist-level classification of skin cancer with deep neural networks*, Nature, **542** (2017), No. 7639, 115–118
- [21] Gigerenzer, G., *Gut Feelings: Short Cuts to Better Decision Making*, Penguin, London, 2008
- [22] Griffiths, T. L., Lucas, C., Williams, J. and Kalish, M. L., *Modeling human function learning with gaussian processes*, Advances in neural information processing systems, **21** (2009), 553–560
- [23] Hinton, G., et al., *Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups*, IEEE Signal Processing Magazine, **29** (2012), No. 6, 82–97
- [24] Hofmann, T., Schoelkopf, B. and Smola, A. J. *Kernel methods in machine learning*, The annals of statistics, **36** (2008), No. 3, 1171–1220
- [25] Holzinger, A., *Human-computer interaction & knowledge discovery (HCI-KDD): What is the benefit of bringing those two fields to work together?*, LNCS **8127** (2013), 319–328
- [26] Holzinger, A., *Interactive machine learning (iml)*. Informatik Spektrum, **39** (2016), No. 1, 64–68
- [27] Holzinger, A., *Interactive machine learning for health informatics: When do we need the human-in-the-loop?*, Brain Informatics, **3** (2016), No. 2, 119–131
- [28] Holzinger, A., Plass, M., Holzinger, K., Crisan, G. C., Pinteau, C-M. and Palade, V., *Towards interactive machine learning (iml): Applying ant colony algorithms to solve the traveling salesman problem with the human-in-the-loop approach*, LNCS, **9817** (2016), 81–95
- [29] Holzinger, A., Plass, M. and Kickmeier-Rust, M. D., *Interactive machine learning (iml): a challenge for game-based approaches*, in *Challenges in Machine Learning: Gaming and Education. NIPS Workshops*, 2016

- [30] Holzinger, A., Plass, M., Kickmeier-Rust, M., Holzinger, K., Crişan, G. C., Pintea, C-M. and Palade, V., *Interactive machine learning: experimental evidence for the human in the algorithmic loop. A case study on Ant Colony Optimization*, Applied Intelligence, (2018) 1–14
- [31] Hund, M., et al., *Visual analytics for concept exploration in subspaces of patient groups: Making sense of complex datasets with the doctor-in-the-loop*, Brain Informatics, (2016), 1–15
- [32] iML Project, HCI-KDD, Medical University Graz <http://hci-kdd.org/project/impl>
- [33] Kieseberg, P., Malle, B., Fruehwirt, P., Weippl, E. and Holzinger, A., *A tamper-proof audit and control system for the doctor in the loop*, Brain Informatics, (2016), 1–11
- [34] Kieseberg, P., Weippl, E. and Holzinger, A., *Trust for the doctor-in-the-loop*, European Research Consortium for Informatics and Mathematics, News: Tackling Big Data in the Life Sciences, **104** (2016), No. 1, 32–33
- [35] Knill, D. C. and Pouget, A., *The Bayesian brain: the role of uncertainty in neural coding and computation*, Trends in neurosciences, **27** (2004), No. 12, 712–719
- [36] Lin, S., *Computer solutions of the traveling salesman problem*, The Bell System Technical Journal, **44** (1965), No. 10, 2245–2269
- [37] Lucas, C. G., Griffiths, T. L., Williams, J. J. and Kalish, M. L., *A rational model of function learning*, Psychonomic Bulletin & Review, **22** (2015), No. 5, 1193–1215
- [38] MacGregor, J. N. and Ormerod, T., *Human performance on the traveling salesman problem*, Perception & Psychophysics, **58** (1996), No. 4, 527–539
- [39] MacGregor, J. N. and Chu, Y., *Human performance on the traveling salesman and related problems: A review*, The Journal of Problem Solving, **3** (2011), No. 2, 119–150
- [40] Malle, B., Kieseberg, P., Weippl, E. and Holzinger, A., *The right to be forgotten: Towards machine learning on perturbed knowledge bases*, LNCS (2016), No. 9817, 251–256
- [41] Malle, B., Kieseberg, P. and Holzinger, A., *Do not disturb? classifier behavior on perturbed datasets*, LNCS (2017), No. 10410, 155–173
- [42] Matei, O., Pop, P. C. and Vălean, H., *Optical character recognition in real environments using neural networks and k-nearest neighbor*, Applied intelligence, **39** (2013), No. 4, 739–748
- [43] Mnih, V., et al., *Human-level control through deep reinforcement learning*, Nature, **518** (2015), No. 7540, 529–533
- [44] Monasson, R., Zecchina, R., Kirkpatrick, S., Selman, B. and Troyansky, L., *Determining computational complexity from characteristic phase transitions*, Nature **400** (1999), No. 6740, 133–137
- [45] Napolitano, F., Raiconi, G., Tagliaferri, R., Ciaramella, A., Staiano, A. and Miele, G., *Clustering and visualization approaches for human cell cycle gene expression data analysis*, International Journal of Approximate Reasoning, **47** (2008), No. 1, 70–84
- [46] Pintea, C-M., Dumitrescu, D. and Pop, P. C., *Combining Heuristics and Modifying Local Information to Guide Ant-based Search*, Carpathian J. Math., **24** (2008) No. 1, 94–103
- [47] Pintea, C-M., Crişan, G-C. and Chira, C., *Hybrid Ant Models with a Transition Policy for Solving a Complex Problem*, Logic Journal of the IGPL, **3** (2012), 560–569
- [48] Pintea, C-M. et al., *A Hybrid Ant-based System for Gate Assignment Problem*, LNCS, **5271** (2008), 273–280
- [49] Pop, C. P., Pintea, C-M. and Sitar, C. P., *An Ant-based Heuristic for the Railway Traveling Salesman Problem*, LNCS, **4448** (2007), 702–711
- [50] Pintea, C-M., Crişan, G-C., Chira, C., *A Hybrid ACO Approach to the Matrix Bandwidth Minimization Problem*, LNCS, **6076** (2010), 405–412
- [51] Pintea, C-M. et al., *A hybrid classical approach to a fixed-charged transportation problem*, LNCS, **7208** (2012), 557–566
- [52] Robert, S., Bttner, S., Rcker, C., and Holzinger, A., *Reasoning under uncertainty: Towards collaborative interactive machine learning*, in (Holzinger, A., Ed.), Machine Learning for Health Informatics: State-of-the-Art and Future Challenges, Springer Cham, (2016), 357–376
- [53] Rooij, I. V., Stege, U. and Schactman, A., *Convex hull and tour crossings in the euclidean traveling salesperson problem: Implications for human performance studies*, Memory & cognition, **31** (2003), No. 2, 215–220
- [54] Russakovsky, O., et al., *Imagenet large scale visual recognition challenge*, International Journal of Computer Vision, **115** (2015), No. 3, 211–252
- [55] Schirner, G., Erdogmus, D., Chowdhury, K. and Padir, T., *The future of human-in-the-loop cyber-physical systems*, Computer, **46** (2013), No. 1, 36–45
- [56] Shahriari, B., Swersky, K., Wang, Z., Adams, R. P. and de Freitas, N., *Taking the human out of the loop: A review of bayesian optimization*, Proceedings of the IEEE, **104** (2016), No. 1, 148–175
- [57] Shyu, C-R. et al., *Assert: A physician-in-the-loop content-based retrieval system for hrct image databases*, Computer Vision and Image Understanding, **75** (1999), No. 12, 111–132
- [58] Silver, D., et al., *Mastering the game of go with deep neural networks and tree search*, Nature, **529** (2016), No. 7587, 484–489
- [59] Sonnenburg, S., Raetsch, G., Schaefer, C. and Schoelkopf, B., *Large scale multiple kernel learning*, Journal of Machine Learning Research, **7** (2006), 1531–1565

- [60] Steyvers, M., Tenenbaum, J. B., Wagenmakers, E-J. and Blum, B., *Inferring causal networks from observations and interventions*, *Cognitive science*, **27** (2003), No. 3, 453–489
- [61] Stoean, C. and Stoean, R., *Support vector machines and evolutionary algorithms for classification*, Springer, 2014
- [62] Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J. and Wojna, Z., *Rethinking the inception architecture for computer vision in IEEE Conference on Computer Vision and Pattern Recognition*, (2016), 2818–2826
- [63] Tenenbaum, J. B., Kemp, C., Griffiths, T. L. and Kemp, C., *Theory-based bayesian models of inductive learning and reasoning*, *Trends in cognitive sciences*, **10** (2006), No. 7, 309–318
- [64] Tenenbaum, J. B., Kemp, C., Griffiths, T. L. and Goodman, N. D., *How to grow a mind: Statistics, structure, and abstraction*, *Science*, **331** (2011), No. 6022, 1279–1285
- [65] Thaker, P., Tenenbaum, J. B. and Gershman, S. J., *Online learning of symbolic concepts*, *Journal of Mathematical Psychology*, **77** (2017), 10–20
- [66] TSP Library: <https://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95/>
- [67] Ventresca, M. and Ombuki, B. M., *Ant colony optimization for job shop scheduling problem*, in *Proceedings of 8th IASTED International Conference on Artificial Intelligence and Soft Computing*, (2004), 451–152
- [68] Wilson, A. G., Dann, C., Lucas, C. and Xing, E. P., *The human kernel*, in *Advances in Neural Information Processing Systems*, NIPS 2015 (Cortes, C. et al. Eds.), **28** (2015), 2836–2844
- [69] Wolpert, D. M., Ghahramani, Z. and Jordan, M. I., *An internal model for sensorimotor integration*, *Science*, **269** (1995), No. 5232, 1880–1882
- [70] Xu, F. and Tenenbaum, J. B., *Word learning as bayesian inference*, *Psychological review*, **114** (2007), No. 2, 245–272

¹HOLZINGER-GROUP, HCI-KDD, MEDICAL UNIVERSITY GRAZ
INSTITUTE FOR MEDICAL INFORMATICS, STATISTICS AND DOCUMENTATION
AUENBRUGGERPLATZ 2/V, A-8036 GRAZ, AUSTRIA

²INSTITUTE OF INTERACTIVE SYSTEMS AND DATA SCIENCE
GRAZ UNIVERSITY OF TECHNOLOGY,
INFFELDGASSE 16C, 8010 GRAZ AUSTRIA
E-mail address: a.holzinger@hci-kdd.org

³VASILE ALECSANDRI UNIVERSITY OF BACĂU,
FACULTY OF SCIENCES,
CALEA MARASESTI 157, 600115 BACĂU, ROMANIA
E-mail address: ceraselacrisan@ub.ro

⁴DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE
TECHNICAL UNIVERSITY OF CLUJ-NAPOCA
NORTH UNIVERSITY CENTRE AT BAI A MARE
VICTORIEI 76, 430122 BAI A MARE, ROMANIA
E-mail address: dr.camelia.pintea@ieee.org

⁵COVENTRY UNIVERSITY
FACULTY OF ENGINEERING, ENVIRONMENT AND COMPUTING
PRIORY STREET CV1 5FB COVENTRY, UNITED KINGDOM
E-mail address: vasile.palade@coventry.ac.uk