



HAL
open science

The Human in Interactive Machine Learning: Analysis and Perspectives for Ambient Intelligence

Kévin Delcourt, Sylvie Trouilhet, Jean-Paul Arcangeli, Françoise Adreit

► **To cite this version:**

Kévin Delcourt, Sylvie Trouilhet, Jean-Paul Arcangeli, Françoise Adreit. The Human in Interactive Machine Learning: Analysis and Perspectives for Ambient Intelligence. *Journal of Artificial Intelligence Research*, 2024, 81, pp.263-305. 10.1613/jair.1.15665 . hal-04756752

HAL Id: hal-04756752

<https://hal.science/hal-04756752v1>

Submitted on 28 Oct 2024

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.



Distributed under a Creative Commons Attribution 4.0 International License

The Human in Interactive Machine Learning: Analysis and Perspectives for Ambient Intelligence

Kévin Delcourt

Sylvie Trouilhet

Jean-Paul Arcangeli

IRIT, Université de Toulouse, UT3, Toulouse, France

KEVIN.DELCOURT@IRIT.FR

SYLVIE.TROUILHET@IRIT.FR

JEAN-PAUL.ARCANGELI@IRIT.FR

Françoise Adreit

IRIT, Université de Toulouse, UT2J, Toulouse, France

FRANCOISE.ADREIT@IRIT.FR

Abstract

As the vision of Ambient Intelligence (AmI) becomes more feasible, the challenge of designing effective and usable human-machine interaction in this context becomes increasingly important. Interactive Machine Learning (IML) offers a set of techniques and tools to involve end-users in the machine learning process, making it possible to build more trustworthy and adaptable ambient systems. In this paper, our focus is on exploring approaches to effectively integrate and assist human users within ML-based AmI systems. Through a survey of key IML-related contributions, we identify principles for designing effective human-AI interaction in AmI applications. We apply them to the case of Opportunistic Composition, which is an approach to achieve AmI, to enhance collaboration between humans and Artificial Intelligence. Our study highlights the need for user-centered and context-aware design, and provides insights into the challenges and opportunities of integrating IML techniques into AmI systems.

1. Introduction

Ambient Intelligence (AmI) aims to provide a personalized physical and software environment that adapts to users' needs and situations (Sadri, 2011; Dunne et al., 2021). Its potential impact is significant, as it can enhance human-environment interaction through the integration of intelligent systems into everyday life. There are various applications of AmI, particularly in the fields of healthcare (Acampora et al., 2013), transportation (Vestin et al., 2004), energy management (Robinson et al., 2015), and smart homes (Makonin et al., 2012).

However, the dynamics, openness, and unpredictability of ambient environments pose significant challenges to the development of effective AmI systems, particularly due to the mobility of devices and users, and the diversity of components in the environment. To overcome these challenges, solutions must take into account the operational context, including user preferences and needs that may vary with time. Opportunistic Composition is an approach to achieve AmI. It revolves around the idea of dynamically and opportunistically constructing complex applications by leveraging existing software components in the environment. The goal of Opportunistic Composition is to enable the seamless integration and collaboration of these components to create adaptive and context-aware systems. It is achieved by the Opportunistic Composition Engine (OCE) (Delcourt et al., 2021), that relies on Machine Learning in interaction with the human user.

Interactive Machine Learning (IML) (Fails & Olsen Jr, 2003) investigates ways to enable humans to teach machine learning algorithms, with a focus on providing tools that are usable by end-users without machine learning backgrounds. As such, IML can provide a valuable set of tools and techniques for addressing the challenges of AmI and enabling more effective human-machine interaction.

The aim of this paper is to survey design solutions and recommendations that enhance the collaboration between a human and a learning machine. Then, we analyze their application to Opportunistic Composition. To achieve this goal, we undertake a critical review (Grant & Booth, 2009) of the existing IML-related literature. By critically analyzing and synthesizing insights from this body of work, we aim to advance Opportunistic Composition while uncovering potential applications beyond the specific context of our project. Indeed, the solutions and recommendations we provide mark the initial steps in addressing the question of designing socially responsible AI solutions (Cheng et al., 2021) that provide a fair, transparent and secure interaction.

The article is structured as follows:

- Section 2 offers a comprehensive overview of the fields of Machine Learning, Interactive Machine Learning, and sets the research questions that this survey aims to address.
- Section 3 provides a thorough analysis of several key IML-related contributions that are relevant to the challenges of AmI. These contributions are analyzed based on the research questions outlined in Section 2, and provide insights into how IML can be used to address the specific challenges of AmI systems.
- Section 4 provides a comprehensive synthesis of the key findings from the previous section.
- Section 5 presents the principles of Opportunistic Composition and the Opportunistic Composition Engine (OCE). Then, it discusses how the findings from Sections 3 and 4 can be applied to Opportunistic Composition to enhance human-AI collaboration within this project.
- Section 6 provides concluding remarks and suggests potential avenues for future research in the areas of Opportunistic Composition, IML, and AmI.

2. Learning in Interaction

In situations where solving complex problems through programming is not feasible, due to a lack of understanding, Machine Learning (ML) is often the go-to method for building a solution. Developing an ML-based solution typically involves engaging an ML expert, as illustrated in Figure 1. This expert is responsible for designing the learning aspect of the solution, such as selecting the appropriate algorithm and tuning the parameters, while working in collaboration with both the end-users and the ML system under development. The users or their representatives provide data to the expert, who then tunes the learning system until it produces satisfactory results. These results are subsequently reviewed by the users, who provide further feedback. This iterative process continues until the ML system, precisely the ML model, is deemed ready for production.

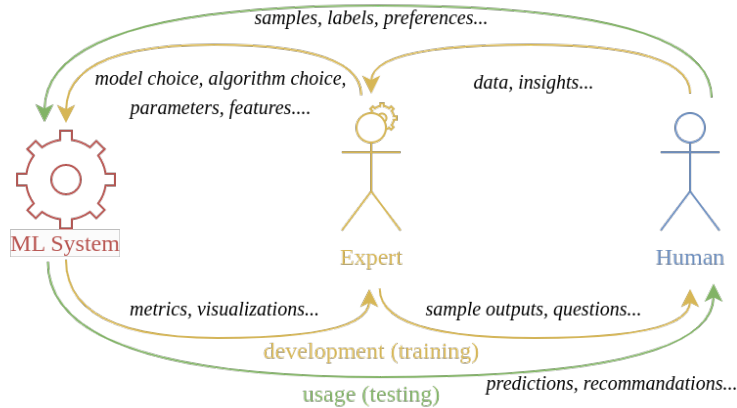


Figure 1: Traditional ML training process, adapted from (Amershi et al., 2014)

However, as noted by Amershi et al. (Amershi et al., 2014), there is a large demand for machine learning applications but a shortage of experts in this field, which can slow down the development process. To address this challenge, Interactive Machine Learning (IML) approaches have been proposed to allow human users, who are the target users of these applications, to contribute to the learning process without requiring an ML expert. These approaches are designed for humans without prior knowledge of machine learning and are categorized under IML. In certain cases, a user representative with knowledge of the user base and their needs may also be considered a human user.

2.1 Interactive Machine Learning

According to (Fails & Olsen Jr, 2003), IML is the field of machine learning that involves human users interacting with a learning algorithm to provide some or all of the data used for learning. The goal of IML is to produce a solution that meets the needs of the end-users of the ML-based solution, with the involvement of humans who typically lack machine learning skills. This human involvement in the learning process enables personalization of the resulting ML system.

The iterative operation of an IML system, as depicted in Figure 2, involves the human providing various parameters, preferences, or any data required by the ML system for its operation. The ML system then presents the results of its learning, such as recommendations or predictions, to the human who evaluates if the presented results are satisfactory with respect to their objectives. The human provides additional inputs and feedback, which are then incorporated by the ML system to update its model. The system can be retested with the updated model, and the process repeats until the human is satisfied with the results. In certain cases, the process may continue indefinitely, while the system is operated, to adapt to new data and changes in user preferences or needs. Compared to the process described in Figure 1, IML feedback loops are faster in the absence of the intermediate expert.

IML applications have been successfully used to build various machine learning solutions such as image classifiers (Carney et al., 2020), text classifiers (Ramos et al., 2020), training

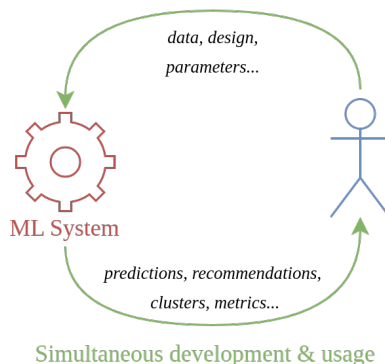


Figure 2: IML training process, adapted from (Amershi et al., 2014)

robots to perform handling tasks (Celemin & Ruiz-del Solar, 2019), and recommendations based on user actions (Aamir & Bhusry, 2015).

Several research areas share similarities with IML and are closely related. AutoML focuses on automating the entire machine learning pipeline, from data pre-processing to model deployment (Karmaker et al., 2021). The primary goal of AutoML is to make machine learning more accessible to non-experts and to improve the efficiency of experts. AutoML is primarily an automated, data-driven approach (Hutter et al., 2019), while IML is a more interactive human-driven approach. While AutoML has the potential to improve the efficiency and effectiveness of machine learning systems, it may not be suitable in cases where data is scarce or not available.

Human-autonomy teaming, also known as human-autonomy collaboration, refers to the collaboration between humans and autonomous agents to achieve a common goal, such as data processing or decision making (O’Neill et al., 2022). This research field tends to focus on the interactions between humans and AI, for instance, how to coordinate the two in a seamless way (Liang et al., 2019). To highlight the use of AI in autonomous agents, Human-autonomy teams are often referred to as Human-AI teams (Zhang et al., 2021a).

Hybrid Intelligence systems are similar to IML, where a human and an AI agent collaborate to solve complex tasks (Dellermann et al., 2019). While the majority of current ML systems do not take into account human feedback after initial training, making them unsuitable for real-world dynamic settings, Hybrid Intelligence theorizes that continuous adaptation of the ML model and continuous human intervention are necessary to achieve results in highly dynamic environments.

Our work focuses on the field of Ambient Intelligence (Sadri, 2011), where dynamics are high, and humans constantly interact with their environment. Adaptation is a fundamental property, and IML is a promising approach for building intelligent systems that can adapt to human needs and preferences in real-time.

2.2 Problems and Research Questions

AmI and IML systems share similarities as both focus on placing the human at the center of the application and adapting the service provided to them. While AmI refers specifically to the assistance of end-users in a pervasive environment, IML refers more broadly to ML systems with the user in the loop. Machine learning is not a requirement for AmI, but it is becoming an increasingly popular method for achieving it (Dunne et al., 2021).

AmI is an ongoing research field that has yet to be fully realized (Stankovic et al., 2021). However, AmI and IML are related research fields, and IML can provide solutions to the challenges of AmI. Therefore, we aim to address these challenges by reviewing IML-related literature, with a focus on issues related to the integration of human users into the ambient system. These include:

- Ensuring a smooth interaction between the human and the ambient system.
- Dynamically adapting and evolving intelligent behavior over time in response to changes in the human’s goals, preferences and environment.
- Ensuring relevant outputs given the inherent dynamics and unpredictability of the environment (networks, behaviors, unexpected events, noise, etc).

Given the importance of effective human integration for successful IML solutions (Amer-shi et al., 2014), we aim to draw insights and guidelines from the IML literature that can be applied to ambient systems. To achieve this goal, we have formulated the following research questions that will guide our analysis of key contributions to the field of IML. The answers to these questions will enable us to identify recommendations in the context of an AmI system, which will be presented in Section 5 based on our discussion in Section 4.

RQ1-Human. What are the human’s role and responsibilities in the loop?

- **RQ11-Tasks.** What tasks should the human perform? To accomplish these tasks, how much machine learning skills are required?

This understanding is important to ensure that the human is not overburdened with tasks that are beyond their skills. It can guide the design of IML systems that take into account the skills and limitations of the human user.

- **RQ12-Workload.** How much and what workload is imposed on the human user? What level of commitment or involvement is expected?

Understanding the workload and level of commitment required from the human user is crucial to ensure that the system is designed to provide the right level of support and assistance to optimize performance and user experience. In the following, to enable a comparison of the workload imposed on human users across different applications, we have defined three distinct levels of workload: light, medium, and heavy. These levels are determined based on our assessment of the time needed to complete tasks with the system, the number of steps involved, and the degree of mental effort required.

RQ2-Learner. How is the human taken into account by the ML system?

- **RQ21-Information.** What information is needed from the human and what is it used for?

The exchanged information can affect the accuracy and reliability of the system. Understanding what information is needed from the human user and how it can be incorporated into the system can help to improve the performance and effectiveness of the IML system. Offering varied and comprehensive methods of interacting with AI systems has been shown to improve both system performance and user experience (Amershi et al., 2014).

- **RQ22-Assistance.** How does the system guide and assist the human?

The provided assistance affects the efficiency, accuracy, and user experience of the system. Providing the right level of guidance and assistance to the human user can help to optimize their performance and ensure that the system operates effectively. To answer this question, we examine how the system delivers context-specific help and feedback to the user as they perform their tasks. Additionally, we assess the system’s ability to offer recommendations and suggestions to the user, enabling them to make informed decisions and complete their tasks with greater efficiency.




Figure 3: Research questions (RQs)

Derived from Figure 2, Figure 3 places the different research questions in relation to the IML process presented in Section 2.1. They cover all aspects of the interaction between the human and the ML system. In fact, these questions go beyond ML and address the relationship between humans and AI more broadly, as discussed, for example, by (Saisubramanian et al., 2022).

3. Relationship Between Humans and Learning Machines in IML

To examine how the IML-related literature addresses the research questions at hand, this paper conducts a critical review of the literature, in accordance to the definition¹ provided by (Grant & Booth, 2009). The choice of this type of review was guided by the fact that our goal is to produce a synthetic model of human-AI collaboration that is applicable to

1. "An effective critical review presents, analyses and synthesizes material from diverse sources. Its product [...] typically manifest in a hypothesis or a model, not an answer. The resultant model may constitute a synthesis of existing models or schools of thought or it may be a completely new interpretation of the existing data."

Opportunistic Composition and generalizable to ML-based AmI systems. This model is synthesized in Section 4.

Our survey is based on a thorough analysis of the IML-related literature. Rather than aiming for an exhaustive coverage of all available literature, our approach involved the identification of 13 relevant scientific contributions that offered valuable insights into our research questions. The selection process for these papers was exploratory and flexible, allowing us to uncover a diverse set of publications that might have been overlooked by a strictly systematic approach. For instance some of the selected papers may have lacked certain keywords such as "Interactive Machine Learning". The selection process was guided by our expertise in the field of AmI, enabling us to identify key papers and influential works that addressed our specific research problems. This expertise-driven approach ensured the inclusion of important contributions in the field and allowed us to focus on the most relevant literature. However, we acknowledge that our approach may introduce biases and impact the generalizability of our findings.

Several state of the art reviews exist, each with a specific scope. Dudley and Kristensson study interactive supervised learning applications (Dudley & Kristensson, 2018) but they do not specifically consider the unique characteristics of Ambient Intelligent applications, such as the dynamic and open nature of the environment. In (Li et al., 2019), work on human-centered reinforcement learning is surveyed with a focus on algorithms that learn from evaluative feedback of a human observer. Zhang, Torabi, Warnell, and Stone are interested in certain specific interaction techniques, including learning from user evaluative feedback (Zhang et al., 2021b). Additionally, several work focus on machine learning visualization (Jiang et al., 2019) and interpretability (Chatzimparmpas et al., 2020).

In this section we analyze the selected contributions. We summarize each of them and then analyze how the authors answer the four research questions in their solutions. These papers cover a variety of works, including IML-based applications (Section 3.1), ML approaches that target a particular issue in relation to the user (Section 3.2), and studies on specific aspects of IML (Section 3.3). Finally, we study a set of generic guidelines for the design of IML applications (Section 3.4). They are a set of best practices and principles, that focus on improving the performance and usability of human-AI interaction.

3.1 Applications

This section regroups five works that propose concrete IML applications, such as data classifiers (Fails & Olsen Jr, 2003; Carney et al., 2020; Flutura et al., 2018; Berg et al., 2019) or recommendation systems (Zheng et al., 2018). For each of them, we analyze and synthesize their answers to the research questions.

3.1.1 IMAGE, SOUND AND POSTURE CLASSIFICATION

Carney et al. introduce Google's Teachable Machine tool² allowing end-users to build a classifier of images, sounds, or human body postures (Carney et al., 2020). The tool is designed for users new to supervised learning, with possible coding skills, since the models can be exported in TensorFlow (Abadi et al., 2016).

2. <https://teachablemachine.withgoogle.com/>

In order to facilitate the users’ task by limiting the training data to provide in order to achieve satisfactory results, a solution based on transfer learning (Weiss et al., 2016) has been implemented: the training of a new model by the user is done on the basis of a versatile basic model, trained upstream by the classifier developers.

We find the following answers to our RQs, as summarized in Figure 4:

RQ11-Tasks. Here the user is in charge of selecting the concepts to be taught (e.g. to classify pictures with or without cats), and collecting the associated examples (e.g. images). They must observe and evaluate the learner’s results themselves, thus providing the test data.

RQ12-Workload. The user is involved in each step of the building of the learned model, from the choice of training data to the evaluation of the learner. This involves them strongly.

RQ21-Information. Human-supplied input, in the form of labeled data, allows the basic model provided by the application to be customized as needed. Besides, it is possible to modify the learner’s advanced settings, e.g. the learning rate.

RQ22-Assistance. The HMI guides the user through the various stages of model production. In addition, it initially hides the advanced settings from the learner, so that the investment to produce a model is as low as possible.




Figure 4: Answers to RQs for Google’s Teachable Machine

3.1.2 PIXEL CLASSIFICATION

Fails and Olsen Jr propose to train a pattern recognition system, using an IML application called Crayons (Fails & Olsen Jr, 2003). Via an interface, the user colors (i.e. labels) areas of the image that they want the learner to classify. The supervised learning process then proposes a classification of the whole image to the user, who can refine their labeling to correct possible errors.

The authors discuss the various requirements associated with IML, in particular that learning should be carried out as quickly as possible so as not to discourage the user. With respect to these requirements, they identify decision trees (Quinlan, 1986) as the best approach to IML, given its short learning time.

Crayons’ answers to our RQs are presented below, as depicted in Figure 5.

RQ11-Tasks. The user must be able to determine on the image the classes to be identified (e.g. the skin) and color a part of it.

RQ12-Workload. The user does not have to precisely color the areas to be classified, thus avoiding mental overload. However, he has to refine the labeling until the results are appropriate.

RQ21-Information. The image provided by the user and the labels they apply while coloring are the source of the supervised learning.

RQ22-Assistance. The user is not particularly accompanied but labeling by coloring makes the user’s task quite natural and accessible to anyone.




Figure 5: Answers to RQs for Crayons (Fails & Olsen Jr, 2003)

3.1.3 MEDICAL IMAGE CLASSIFICATION

Ilastik is an interactive supervised learning tool for image classification in the medical field (Berg et al., 2019). It allows a medical expert to produce an image classifier. Compared to Crayons (Section 3.1.2), it takes into account data up to 5 dimensions (space, time, and number of channels). The tool also offers the user a choice of seven image classification modes (e.g. pixel classification or object classification), which involve different interactions with the user. For example, the way in which the examples provided by the user are labeled (coloring, clicking on shapes to be detected, etc.) depends on the chosen algorithm.

Below are Ilastik’s answers to our RQs, which are summarized in Figure 6.

RQ11-Tasks. The user must have skills in the medical field. In addition, they must choose which classification mode to use, according to their objective. They must then provide and annotate representative examples.

RQ12-Workload. The diversity of classification modes, which involves different actions on the part of the user, means that they must be strongly invested in the application in order to master it.

RQ21-Information. The annotated examples provided by the user are the only source of training.

RQ22-Assistance. Except for the documentation available online, the user has little support to accomplish their task. A library of pre-trained models is proposed, but it is not possible to alter these models to adapt them to a particular task.

3.1.4 ACTIVITY DETECTION

Flutura et al. have developed Drinkwatch, a supervised learning application embedded in a connected watch that allows users to track their beverage consumption (Flutura et al.,




Figure 6: Answers to RQs for Ilastik (Berg et al., 2019)

2018). To do this, the application learns the consumption habits of the user via the sensors of the watch.

Figure 7 summarizes Drinkwatch’s answers to our RQs, which are presented below.

RQ11-Tasks. Throughout its activity, the user is called upon to validate or invalidate certain activities detected by the watch, i.e. they collaborate to detect false positives. They should also report cases where activity has not been detected, i.e. indicate false negatives.

RQ12-Workload. The user is solicited only when the application has a low degree of confidence in its classification, to avoid overloading the human. However, they remain involved since they have to be vigilant to the false negatives and positives of the application.

RQ21-Information. The learning is both based on the movements detected by the watch for pattern recognition and the explicit feedback from the user.

RQ22-Assistance. In addition to the fact that the application solicits as little human effort as possible, sound notifications signal the detection of activity by Drinkwatch in order to facilitate the user’s work.



Figure 7: Answers to RQs for Drinkwatch (Flutura et al., 2018)

3.1.5 NEWS RECOMMENDATIONS

In (Zheng et al., 2018), authors propose a deep reinforcement learning framework for news recommendations, called Deep Reinforcement News (DRN). The challenges in this domain are twofold: first the news articles become quickly outdated, and second the users have their own goals and preferences that evolve over time. The learning agent personalizes its recommendations by taking into account implicit feedback from the user. A single learning agent processes feedback and recommendations for all users. Other than the news chosen or not by the user, additional data such as the time taken to come back to the service, e.g.

the user’s general interest in the application, serve as rewards to the reinforcement learning agent. Thus the goal is to maximize both long term and immediate rewards.

Additionally, the authors deal with the challenge of the dynamics of a user’s goals and taste by implementing an effective exploration strategy, that recommends novel items in accordance to the known tastes of the user.

The following presents DRN’s answers to our RQs, as outlined in Figure 8.

RQ11-Tasks. Apart from browsing the news recommendation service and reading the articles, the user has no other task to perform.

RQ12-Workload. The user is only required to browse the service for the learning to take place. There is no additional workload on the user.

RQ21-Information. The learning agent takes into account several features from the news and the user. The detail of the usage of the application by the user (clicks, time spent on the service, time to return, etc.) is implicitly gathered and exploited to personalize the recommendations.

RQ22-Assistance. The recommendation service only allows implicit interaction with the learning agent, and thus does not need to assist the user in additional ways.



Figure 8: Answers to RQs for Deep Reinforcement News (Zheng et al., 2018)

3.2 Approaches

This section presents four approaches that particularize IML (Ramos et al., 2020) or adapt ML-based solutions (Kessler Faulkner & Thomaz, 2021; Akrouf et al., 2014; Christiano et al., 2017) in relation to the human in the loop, its role or involvement. For each of the applications developed within those approaches, we analyze and synthesize their answers to the research questions.

3.2.1 INTERACTIVE MACHINE TEACHING

Defined in (Ramos et al., 2020), Interactive Machine Teaching is an approach to IML in which the human plays the role of a teacher, and must teach a task to the machine. The notion of teaching includes the choice of information at the source of learning, and the evaluation of the learner’s performance. The underlying assumption that the authors make is that humans acquire pedagogical skills more easily than machine learning skills, as these skills are more prevalent in the general public (Wall et al., 2019).

The concepts and processes studied in Interactive Machine Teaching are applicable to any learning paradigm. Nevertheless, the authors present a demonstration application,

called PICL (Ramos et al., 2020) (formerly MATE (Wall et al., 2019)), that applies Interactive Machine Teaching principles to supervised learning by allowing users to teach text classification tasks. As shown in Figure 9, we consider PICL, not Interactive Machine Teaching, as the reference for studying the answers to the RQs.

RQ11-Tasks. The human, playing here the role of a teacher, must plan a curriculum and then update it according to the learner’s results. A curriculum refers to the data (examples, labels) used by the learner. The skills required are mostly of a pedagogical nature.

RQ12-Workload. Involving the human in selecting relevant examples or teaching concepts moderately engages them in the process. In addition, they must judge whether the learner’s results are satisfactory.

RQ21-Information. The labels and concepts provided by the human are the source of supervised learning.

RQ22-Assistance. The interface of the presented tools (Ramos et al., 2020) assists the human by allowing them to effectively visualize the learner’s results. A study of the behavior of supervised learning experts on PICL/MATE (Wall et al., 2019) has also identified good practices for automatic teaching, implemented in the form of notifications.




Figure 9: Answers to RQs for Microsoft’s PICL

3.2.2 LEARNING FROM INATTENTIVE HUMANS

Kessler Faulkner and Thomaz are interested in the problem of inattention of human users (Kessler Faulkner & Thomaz, 2021). To do so, they consider a use case where a robot learns by reinforcement how to perform specific actions (e.g. sorting shapes) with the assistance of a human supervisor. The human observes the robot acting and can positively reward the robot’s actions that are appropriate for them. Then the robot has the choice between exploiting actions that it knows be satisfactory for the human, and exploring new actions with uncertain reward.

Thus, by assuming that the learning agent knows how to detect the absence of human attention, the authors propose that the agent favors exploitation when the human is inattentive, and exploration in the opposite case. Rather than demanding the human’s attention, which can lead to a negative user experience, the robot adapts to the available attention. The answers to our RQs are outlined below and summarized in Figure 10.

RQ11-Tasks. The human must monitor the learner’s actions and, via a simple interface, reward those that are appropriate .

RQ12-Workload. Although learner monitoring tends to be highly human-intensive, the fact that a drop in attention does not adversely affect the learner is beneficial to the user experience.

RQ21-Information. The rewards given by the user are those of reinforcement learning that the learner seeks to maximize.

RQ22-Assistance. Beside the fact that the robot avoids needlessly interrupting the human, it does not provide significant assistance to them.




Figure 10: Answers to RQs for the Reinforcement Learning approach (Kessler Faulkner & Thomaz, 2021)

3.2.3 LEARNING DESPITE ERRONEOUS FEEDBACK

An approach to interactive learning by reinforcement called programming by feedback is presented in (Akroun et al., 2014). The learner’s actions are grouped into short sequences, that the human supervisor must compare in order to generate feedback. For instance, in the cartpole problem (Geva & Sitte, 1993), the learner has to control the movements of the cart and the human must distinguish which movements improve the stability of the pole. A particularity of this approach is that the learner estimates the human’s error rate. Here, supervisor opinions that seem to be the opposite of previous opinions are ignored.

Below, you will find the answers to our RQs, as depicted in Figure 11.

RQ11-Tasks. After each sequence of actions performed by the learner, the human must compare the last sequence with the best sequence among the previous ones, and choose the one that best corresponds to the task at hand. They must therefore be able to evaluate the learner’s behavior globally and not action per action.

RQ12-Workload. The human must study each sequence of actions of the learner, which requires a strong involvement.

RQ21-Information. The rewards from the comparisons produced by the human allow the learner to estimate a reward function that they try to maximize through their actions.

RQ22-Assistance. Here, the human is little assisted in the evaluation of the sequences of actions.




Figure 11: Answers to RQs for the Reinforcement Learning approach (Akrouer et al., 2014)

3.2.4 DEEP REINFORCEMENT LEARNING

Christiano et al. have designed a deep reinforcement learning algorithm based on human feedback (Christiano et al., 2017). To create a neural network based on human feedback, the learner submits short sequences of actions to the human for comparison. The human must choose between two sequences the best one given the task to learn (e.g. perform a backflip). They keep doing so until there are no more comparisons to make, or they decide the results are satisfactory. These comparisons are the source of the learning of a reward function that the algorithm is constantly trying to optimize.

The authors have tested their solution on several use cases via the OpenAI Gym platform (Brockman et al., 2016). In cases where the reward functions are known, the solution can be compared with a classic reinforcement learning algorithm and achieves better results in several cases. However, only one human user participated in the experiments, which limits the scope of the conclusions.

Figure 12 provides a summary of the answers to our RQs, which are presented below.

RQ11-Tasks. Here the user has to compare two sequences of actions and decide on the best one with respect to their objective.

RQ12-Workload. The proposed solution requires human input for less than 1% of the learner’s actions: only the actions for which the learner is not sure are submitted for comparison. The burden on the human is therefore reduced.

RQ21-Information. The feedback from the human allows the algorithm’s reward function to be updated.

RQ22-Assistance. The interface offers keyboard shortcuts to facilitate human interaction.

3.3 Studies

This section presents studies on specific aspects of IML, that respectively focus on the question of trust in IML systems (Honeycutt et al., 2020), the effect of human interaction on the performance of the intelligent system (Holzinger et al., 2019), and the effect of HMI transparency on the user experience (Schnabel et al., 2020). In the three cases, we analyze and synthesize the answers to the research questions, taking as a reference the best version of the experimented application.




Figure 12: Answers to RQs for the Deep Reinforcement Learning approach (Christiano et al., 2017)

3.3.1 TRUST IN IML

Honeycutt et al. conducted a study that addresses the issue of user trust in machine learning applications (Honeycutt et al., 2020). It is based on results in psychology (Van den Bos et al., 1996) which show that the confidence of an individual towards a human decisional group increases if an opinion expressed by the individual is taken into account by the group. Conversely, the individual will have less confidence if his opinion is ignored. The objective of this study is to recover these results by replacing the human decision group with an automatic learner, in this case an interactive supervised learning application for face recognition in images.

The online experiment measured the confidence of human users towards this application with or without interaction on the one hand, and with increasing, constant or decreasing learner performance on the other hand. In practice, the participants had to check and correct the learner's errors.

The results unexpectedly have shown that in general the interacting group has less confidence in the system than the non-interacting group. The explanation put forward by the authors is that the interacting group spent more time focused on the system errors to correct them.

As illustrated in Figure 13, we consider the face recognition application that supported the experiment to look at the research questions.

RQ11-Tasks. The expected skills for the users are only related to the task at hand: face recognition. The human must recognize and correct the mistakes made by the system, e.g. faces that are not detected or falsely detected.

RQ12-Workload. The user is only involved in the tasks of checking and correcting the learner's output.

RQ21-Information. The simulated application takes into account feedback from users on its errors in order to refine its model.

RQ22-Assistance. Users are voluntarily given little guidance in order to measure their subjective appreciation of the application's performance, which is a measure proportional to the confidence they feel in the system (Yin et al., 2019).




Figure 13: Answers to RQs for the face classification application (Honeycutt et al., 2020)

3.3.2 INTERACTIVE LEARNING FOR OPTIMIZATION

Holzinger et al. present an experiment comparing human-assisted and non-human-assisted learning approaches in the context of traveling salesman problem solving (Holzinger et al., 2019). On the basis of a multi-agent Ant Colony Optimization system (Dorigo et al., 2006), two solutions, with or without human participation, are compared. In the solution that involves human interaction, the interface is designed as a snake-like game, in which the human player guides the snake through the points of the traveling salesman problem. By doing so, the path taken by the human player serves as guidance to the ACO algorithm, helping it to find a better solution.

In the presented study, the addition of the human assistant to the multi-agent system improves the results of the optimization algorithm. Figure 14 summarizes the answers to our RQs, which are detailed below.

RQ11-Tasks. The user acts here in an implicit way on the learning. They are asked to play and guide the snake without being informed of the multi-agent system they are influencing.

RQ12-Workload. They are weakly involved, mainly because they are not informed of the learning, nor of its results. Also, the gamification of the interaction reduces the workload.

RQ21-Information. The actions of the human in the game have the effect of helping the multi-agent system.

RQ22-Assistance. Here, the human is not accompanied, but he does not have to be since his contribution is not explicit: he is not informed of the algorithm in the background.




Figure 14: Answers to RQs for the interactive Ant Colony Optimization system (Holzinger et al., 2019)

3.3.3 TRANSPARENCY IN IML

In (Schnabel et al., 2020), authors are interested in the visibility that the user has on the effect of their choice or action on the learner’s learning. They experimented with different interfaces for an article recommendation application where users act by selecting or not articles to read. One of the experimentation’s parameters was the immediacy of the update of the recommendation list. It was either previewed, or displayed right after a user action or delayed. For the previewed version, which was preferred during the study, the consequences of selecting an article to read are highlighted before the action is performed. That is, when the mouse hovers on an article, the changes that its selection would bring to the list of recommendations are highlighted.

The objective of the study was to measure the impact of this visibility on the efficiency of the users, measured in number of good recommendations accepted by the human in a given time. The study also looks at the perceived user experience. They concluded that having a previewing mechanism was favored by people over other more delayed updates common in most recommendation systems.

To analyze the answers to the research questions outlined in Figure 15, we consider the version of the application that the users have preferred: where the consequences of the user’s actions are previewed. This pre-visualization of the impact of human actions is an assistance that was highly appreciated during the study. It allowed the group involved to be more effective in selecting items.

RQ11-Tasks. The user has to select articles to read among recommendations. They must judge if each news article and the effect of their selection are relevant to their preferences.

RQ12-Workload. The level of human involvement is low considering the task at hand. In the experiment, a time constraint was introduced, which adds a certain mental load to the humans.

RQ21-Information. The actions of the user to add an item to their reading list are the basis of the personalization of the recommendation system.

RQ22-Assistance. In this contribution, the focus is on providing assistance to the humans in the loop. It is greatly effective here.



Figure 15: Answers to RQs for the news recommendation application (Schnabel et al., 2020)

3.4 General Guidelines for Human-AI Applications

Amershi et al. proposed a comprehensive set of 18 guidelines to help designers avoid negative user experiences due to a poorly designed interaction (Amershi et al., 2019). These guidelines are not specific to Interactive Machine Learning applications but are intended for the design of Human-AI interactions in general. The complete guidelines are listed in Table 1.

Guidelines for Human-AI interaction	
Initially	<p>G1. Make clear what the system can do. Help the user understand what the AI system is capable of doing.</p> <p>G2. Make clear how well the system can do what it can do. Help the user understand how often the AI system may make mistakes.</p> <p>GA. Make clear how the system works. Help the user understand the inner workings of the AI system.</p>
During interaction	<p>G3. Time services based on context. Time when to act or interrupt based on the user’s current task and environment.</p> <p>G4. Show contextually relevant information. Display information relevant to the user’s current task and environment.</p> <p>GB. Provide interpretable output to the user. Help the user understand the AI system’s productions.</p> <p>G5. Match relevant social norms. Ensure the experience is delivered in a way that users would expect, given their social and cultural context.</p> <p>G6. Mitigate social biases. Ensure the AI system’s language and behaviors do not reinforce undesirable and unfair stereotypes and biases.</p>
When wrong	<p>G7. Support efficient invocation. Make it easy to invoke or request the AI system’s services when needed.</p> <p>G8. Support efficient dismissal. Make it easy to dismiss or ignore undesired AI system services.</p> <p>G9. Support efficient correction. Make it easy to edit, refine, or recover when the AI system is wrong.</p> <p>G10. Scope services when in doubt. Engage in disambiguation or gracefully degrade the AI system’s services when uncertain about a user’s goals.</p> <p>G11. Make clear why the system did what it did. Enable the user to access an explanation of why the AI system behaved as it did.</p>
Over time	<p>G12. Remember recent interactions. Maintain short term memory and allow the user to make efficient references to that memory.</p> <p>G13. Learn from user behavior. Personalize the user’s experience by learning from their actions over time.</p> <p>G14. Update and adapt cautiously. Limit disruptive changes when updating and adapting the AI system’s behaviors.</p> <p>G15. Encourage granular feedback. Enable the user to provide feedback indicating their preferences during regular interaction with the AI system.</p> <p>G16. Convey the consequences of user actions. Immediately update or convey how user actions will impact future behaviors of the AI system.</p> <p>G17. Provide global controls. Allow the user to globally customize what the AI system monitors and how it behaves.</p> <p>G18. Notify user about changes. Inform the user when the AI system adds or updates its capabilities.</p>

Table 1: Guidelines for Human-AI interaction, from (Amershi et al., 2019). GA and GB were added for our analysis in Section 5.2

Direction of the interaction	Guidelines
Mainly from Human to AI	G7, G8, G9, G12, G15, G17
Mainly from AI to Human	G1, G2, GA, G3, G4, GB, G10, G11, G13, G14, G16, G18

Table 2: Amershi et al.’s guidelines categorized according to the direction of the interaction they pertain to

To produce these guidelines, the authors began with a review of the literature on the subject over the past 20 years. This review yielded 168 specific recommendations, which were consolidated, filtered and refined to 18 proposals. This refinement process consisted of a survey of HMI practitioners and experts, who assessed the relevance of these guidelines in various AI applications.

These general guidelines provide relevant guidance to HMI professionals. This includes, for example, clearly explaining the consequences of the human’s actions on the learner, or conversely explaining how a decision was made by the learning machine. (Amershi et al., 2019) categorized those guidelines based on the different stages of interaction with the AI system.

- Initially: Guidelines related to the initial setup of the Human-AI interaction, including transparency and trust-building.
- During interaction: Guidelines related to how the AI system should adapt to the cultural and environmental human context.
- When wrong: Guidelines related to how the AI system should behave when it makes mistakes or errors. Human users need to understand and intervene when that happens.
- Over time: Guidelines related to how the AI system should evolve and improve over time, including providing opportunities for users to provide feedback to the AI system.

In Table 2, we have organized the guidelines according to the direction of the interaction they pertain to. This categorization helps in understanding the specific areas where Human-AI interactions may be improved. Some guidelines mainly focus on the actions that humans can take towards the AI. For example, guidelines G7, G8, and G9 state that humans should be able to easily invoke, dismiss, and correct the AI system. Other guidelines mainly focus on the different actions the AI can take towards humans. Guidelines such as G10 and G18 suggest that the AI system should ask for human help when it is uncertain of its decision and inform humans of any changes in its capabilities. Finally, G5 and G6 are related to the overall conception and behavior of the AI and are not included in this categorization. They recommend that the AI system should be socially acceptable and avoid hurtful biases and stereotypes.

The guidelines for human-AI interaction proposed by Amershi et al. provide a general framework that can be adapted to different IML applications, effectively answering all four research questions in a broad way, as summarized in Figure 16. As such, we use these guidelines to analyze our solution for ambient intelligence, in Section 5.2.

RQ11-Tasks. Some guidelines suggest types of inputs for the human to perform, such as providing granular feedback to the IML system. The designers should also pay special attention to the invocation, dismissal and correction of the learner.

RQ12-Workload. The guidelines aim to limit human workload by providing only relevant information and interrupting the current task only when necessary. Following these principles, an IML application should not burden the human with unnecessary workload. Other guidelines, however, tend to increase the workload for humans, such as providing feedback at multiple points of the interaction or asking for their assistance in dealing with uncertainty. Although this added load is necessary and beneficial for the interaction, overall it can be quite burdensome.

RQ21-Information. Other guidelines suggest providing additional feedback to the ML algorithm or asking for guidance from the human when the model uncertainty is too high. These additional informations should help to tailor the IML application to the human.

RQ22-Assistance. The major part of the guidelines aim to provide assistance to the human in their interactions with the IML system. This includes making clear what the IML system can do and how well it can do it, building a trusty relationship between the human and the ML algorithm. This also includes working towards explainable AI (Arrieta et al., 2020) to improve human understanding of the learner’s functioning.




Figure 16: Answers to RQs for Microsoft Research’s Human-AI guidelines

4. Analysis of the Contributions and Answers to Research Questions

In this section, we present a comprehensive analysis of the responses to the four research questions, listed in Section 2.2, from the various contributions listed in Section 3. To assist us in our analysis, we have created two tables. Table 3 outlines the answers provided by each contribution, while Table 4 summarizes the level of response to each of the four research questions. In the following we provide a detailed analysis for each research question based on the findings from the various contributions. It is worth noting that the ratings presented in Table 4 for the guidelines of Amershi et al. (Amershi et al., 2019) are based on an hypothetical application implementing all guidelines.

4.1 RQ11-Tasks

Regarding Interactive Machine Learning applications, there is a diverse range of tasks that require human involvement, as shown in column RQ11 of Table 3. Regarding all the contributions, some entail complex tasks, such as providing samples and labels to train the

Contribution	RQ11	RQ12	RQ21	RQ22
Carney et al., 2020	Selection of concept, samples, parameters	Heavy	Labeled images, parameters	Guiding HMI
Fails and Olsen Jr, 2003	Coloring of portions of an image	Medium	Base image, labeled pixels	Request for a natural task
Berg et al., 2019	Choice of workflow, parameters, samples	Heavy	Algorithm, labeled data, parameters	No assistance
Flutura et al., 2018	Reporting of false positive/negatives	Medium	Body movements, explicit feedback	Targeted sound notifications
Zheng et al., 2018	Browsing of news recommendations	Light	Usage data	No assistance
Ramos et al., 2020	Selection of examples and labels	Medium	Labeled textual data	Results' visualization, expert patterns
Kessler Faulkner and Thomaz, 2021	Monitoring and rewarding	Light	Positive rewards	No assistance
Akrour et al., 2014	Comparison of every action sequences	Heavy	Reinforcement feedback	No assistance
Christiano et al., 2017	Comparison of some action sequences	Medium	Reinforcement feedback	Shortcuts
Honeycutt et al., 2020	Assessment of face recognition	Light	Performance feedback	No assistance
Holzinger et al., 2019	Interaction with a snake-like game	Light	Guidance of a MAS	No assistance
Schnabel et al., 2020	Selection of news to read	Light	Usage data	Pre-visualization of the recommendation update
Amershi et al., 2019	Providing granular feedback	Light	Additional feedback, guidance from the human	Building trust, explainability

Table 3: Summary of the answers to the research questions

ML model (Carney et al., 2020; Fails & Olsen Jr, 2003; Ramos et al., 2020). On the other hand, others involve less and even no explicitly ML-oriented tasks, such as playing a game (Holzinger et al., 2019) or selecting articles to read (Zheng et al., 2018; Schnabel et al., 2020).

In addition, column RQ11 of Table 4 rates these tasks as requiring more ML skill (rated -), less ML skill (rated +) or no ML skill at all (rated ++). In general, the presented contributions require few learning skills (+ or ++). However, those rated + still require the human to manipulate ML concepts, such as labels in supervised learning (Ramos et al., 2020; Carney et al., 2020; Fails & Olsen Jr, 2003; Flutura et al., 2018) and action sequences in reinforcement learning (Akrour et al., 2014; Christiano et al., 2017). Furthermore, Ilastik (Berg et al., 2019) requires the human to be familiar with different image recognition al-

Contribution	RQ11	RQ12	RQ21	RQ22
Carney et al., 2020	+	-	++	+
Fails and Olsen Jr, 2003	+	+	-	+
Berg et al., 2019	-	-	++	-
Flutura et al., 2018	+	+	+	+
Zheng et al., 2018	+	++	+	-
Ramos et al., 2020	+	+	+	++
Kessler Faulkner and Thomaz, 2021	++	++	+	-
Akrour et al., 2014	+	-	-	-
Christiano et al., 2017	+	+	-	+
Honeycutt et al., 2020	++	++	-	-
Holzinger et al., 2019	++	++	-	-
Schnabel et al., 2020	+	++	+	++
Amershi et al., 2019	+	++	+	++

Table 4: Summary of the levels of response to the different research questions (the scale ”-” to ”++” indicates the level of response to a question)

gorithms, which is a task closely tied to ML and is thus rated -. Contributions rated ++ involve tasks that are exclusively related to the domain of the application, such as recognizing faces (Honeycutt et al., 2020) or pathfinding (Holzinger et al., 2019).

In summary, a trend among the studied contributions is to initially hide the ML concepts and gradually make them available to the human to understand without overwhelming them. With the exception of (Holzinger et al., 2019), the ML aspects of the interaction are never entirely concealed and are instead discreetly communicated to the human.

4.2 RQ12-Workload

We assessed the workload imposed on the human in the IML applications listed in Table 3 and represented it on a Light-Medium-Heavy scale in column RQ12. The same information is also presented in column RQ12 of Table 4 using a scale ranging from - for heavy to ++ for light. We opted to include the same data on both tables since summarizing each workload in a few words proved challenging in Table 3.

Our analysis indicates that most of the listed contributions impose a medium or light workload on the human, implying that users are less likely to be overwhelmed. This finding perhaps can be explained by the fact that these applications require relatively simple tasks from users, such as pointing out errors made by the machine learner (Honeycutt et al., 2020) or indirectly contributing to learning (Holzinger et al., 2019; Zheng et al., 2018).

Contributions rated heavy require a high level of involvement in the inner workings of the ML applications, which can be a tedious or complex task. In (Akrour et al., 2014), the feedback process requires the human to provide feedback for every action of the learner, which can be a significant burden and demands a considerable amount of effort. Besides, Ilastik (Berg et al., 2019) requires users to choose the algorithm and parameters, which is a complex task. Teachable Machine (Carney et al., 2020) offers the option of including tasks related to the inner workings of the ML system, such as parameter tuning, which may increase the workload in certain cases, but these tasks are not mandatory.

4.3 RQ21-Information

The information provided by the human is crucial for the learner to function, and the way it is used varies between contributions. In Table 3, column RQ21 summarizes what kind of human data fuels the ML system, while in column RQ21 of Table 4, the richness of the data is indicated. Low-dimensional, simpler data is rated - while higher dimensional data is rated + or ++.

Contributions rated ++ offer richer interactions, allowing humans to steer the ML process more effectively towards their preferences. For example, Teachable Machine (Carney et al., 2020) offers deeper access to the learner’s settings, but requires more learning skills from the human. Some implicit sources of interaction, such as mouse hovering (Schnabel et al., 2020) or attention detection (Kessler Faulkner & Thomaz, 2021), allow the human to indirectly influence the learner. These methods provide non-intrusive ways to support the learning process and are accessible to non-specialist end-users.

On the other hand, contributions rated - require simple and low-dimensional data from the human. For instance, interactive reinforcement learning approaches (Akrouf et al., 2014; Christiano et al., 2017) only require a binary reinforcement signal from the human. Although limited interaction between the learner and the human is not necessarily a bad thing and may be enough, research shows that providing rich and diverse ways of interacting with AI systems often leads to better performance and user experience (Amershi et al., 2014).

4.4 RQ22-Assistance

The assistance provided to the human is summarized in column RQ22 of Table 3 and its level is rated on a scale to - to ++ in column RQ22 of Table 4.

In the majority of applications, we found no significant level of assistance provided to the humans, hence rated -. Some applications, like the face recognition system (Honeycutt et al., 2020) and the snake-like game (Holzinger et al., 2019), were designed without any assistance to the humans, likely due to their research nature.

However, contributions rated + and ++ offer various techniques to alleviate the workload for humans. Interactive Machine Teaching (Wall et al., 2019), for example, provides notifications and advice based on expert knowledge in machine learning. In the recommendation application (Schnabel et al., 2020), previsualization techniques allow humans to see the consequences of their actions on the machine learning process, thereby improving their experience. Additionally, guidelines from (Amershi et al., 2019) can be a useful tool for providing assistance in any application because they offer a comprehensive framework for designing human-AI collaborative systems. They address the assistance needs we identified in the literature by advocating for various measures, such as providing comprehensible explanations of the ML system’s purpose or performance, among other things.

The need for new techniques and further research for all these RQs is clear, and could be beneficial to many of the applications studied. In the next section, we examine how the previously identified design solutions can be implemented in the case of opportunistic composition. These design solutions may be transferable to other AmI applications, offering

potential solutions to the challenges related to the interaction between humans and ambient systems that we have identified in Section 2.2.

5. IML in OCE, a Human-Centered Ambient Intelligent Application

As defined in Section 1, AmI aims to provide a personalized and adaptive physical and software environment to meet the needs of the human user (Sadri, 2011). This section applies the findings from Section 3 and Section 4 regarding IML and the relationship between humans and learning machines to the case of Ambient Intelligence (AmI) and specifically Opportunistic Composition. To achieve this, we utilized as general framework the HAX toolkit³, a Microsoft Research tool based on the guidelines presented in (Amershi et al., 2019) (see Table 1), and applied it to the Opportunistic Composition Engine (OCE).

We begin by introducing OCE and its functioning, along with OCE’s answers to the four RQs. Subsequently, following the guidelines of Amershi et al., as well as our own, we investigate how each of these guidelines can be applied in the case of AmI and opportunistic software composition. This investigation is informed by relevant elements seen in Sections 3 et 4.

5.1 Opportunistic Software Composition

Our research project aims at designing and developing a solution that integrates the human user in the decision and learning loop and makes new functionalities emerge in an ambient environment. It is about building the right applications for a human being at the right time, depending on the ambient context, without them having to explicitly ask for these applications. Anticipating and adapting to user needs relies on IML.

It can be noted that while different approaches aim to maximize non-obtrusiveness and transparency for the user, leading to ambient environments that reconfigure themselves automatically, OCE arbitrates in favor of the human’s authority over the system: with OCE, the human ultimately controls the modifications of their environment and can even act directly on it. So the need for smooth interaction between them and the system is all the more important.

An opportunistic composition model has been designed and a software prototype has been developed (Younes et al., 2020; Trouilhet et al., 2021). The Opportunistic Composition Engine (OCE) builds assemblies of software components that are present in the ambient environment at the time. It uses reinforcement learning (Sutton & Barto, 2018), distributed within a multi-agent system (MAS), which takes advantage of user feedback to build knowledge about the user’s preferred components and functionalities. A demonstration of the prototype based on a use case is presented in (Delcourt et al., 2021), and a video of the demonstration is available⁴.

The basic concepts that the human interacting with OCE needs to understand are highlighted in Figure 17. Software components provide services and require services in order to operate. For instance the Display component from Figure 17 provides a service that can set a text to display, while the Temperature Sensor component requires a service

3. <https://www.microsoft.com/en-us/haxtoolkit/workbook/>

4. <https://www.irit.fr/~Sylvie.Trouilhet/wp-content/uploads/sites/90/2022/12/outletSeeking.mp4>

to send its measures to. The two services are connected, symbolized by the dotted line between the two connectors: the temperature is sent to the display. An assembly is a set of components linked together through a set of connections, which represents a usable ambient application. In this example, the measured temperature is displayed and also converted to speech, which is played out loud.




Figure 17: Component diagram with assembly, component, service, and connection highlighted

The iterative process of opportunistic composition is described in Figure 18. OCE probes the human’s ambient environment; its MAS then assembles an application that is proposed to the human for validation; finally, the assembly visualization interface, ICE (Interactive Control Environment) (Trouilhet et al., 2021), allows the human to accept, modify, or reject the OCE proposal. This human action constitutes the feedback that is the source of learning for the OCE agents.

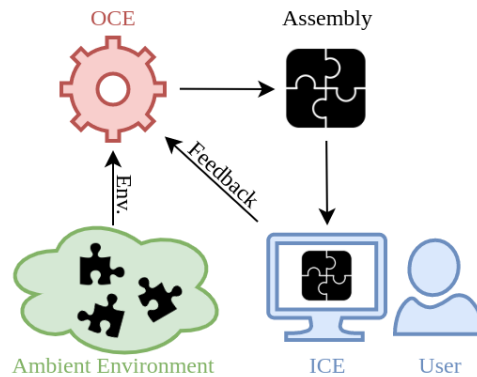


Figure 18: Opportunistic Software Composition

As outlined in Figure 19, the following analyzes how OCE answers to the four research questions. In addition, Table 5 employs the same scale as Table 4 to rate OCE’s level of response.

RQ11-Tasks. To ensure the assembly suits their needs, the human has the ability to accept or reject it, and can also edit connections within the assembly before accepting. This requires an understanding of the function of each software component and the overall purpose of the assembly. Although knowledge of software components and their functionalities



Figure 19: Answers to RQs for the Opportunistic Composition Engine

Contribution	RQ11	RQ12	RQ21	RQ22
Younes et al., 2020	++	-	-	+

Table 5: Summary of the levels of response to the different research questions for OCE (the scale ”-” to ”++” indicates the level of response to a question)

is necessary, no expertise in machine learning is required to interact with OCE, and for that reason we rate it ++.

RQ12-Workload. Here, the human is heavily involved. Accepting, rejecting, and specifically modifying an assembly requires a significant level of engagement and skills from the human. Additionally, in the current OCE prototype, every change in the user’s environment triggers the generation of a new assembly proposal, which must be carefully evaluated by the human, and so we rate it -.

RQ21-Information. The learning mechanism of OCE extracts data from the human feedback, which consists of accepting, rejecting, or modifying an assembly. While this interaction is essential for the system’s learning process, it is limited in scope and is thus rated -. Furthermore, the human user is not aware that their actions have an impact on OCE’s learning mechanism.

RQ22-Assistance. In order to improve the user experience of OCE, (Trouilhet et al., 2021) proposed model transformations to generate various representation of the assemblies, which are illustrated in Figure 20. These range from UML⁵ diagrams to more semantically accurate and user-friendly graphical descriptions. This work could be extended to provide improved assistance, such as generating textual descriptions of assemblies that are easily understandable. For this reason we rate OCE’s assistance +.

5.2 Enhancing Human-AI Collaboration in OCE

In this section, our focus is on improving OCE’s response to the RQs, which involves enhancing human-AI collaboration in the context of an AmI application. To accomplish this goal, we identify design solutions from Section 3 and explore how they could be applied to AmI and OCE. We have organized this study using the guidelines of Amershi et al. (Amershi et al., 2019), which offer a comprehensive framework for answering all of our

5. Unified Modeling Language : <http://www.omg.org/spec/UML/>

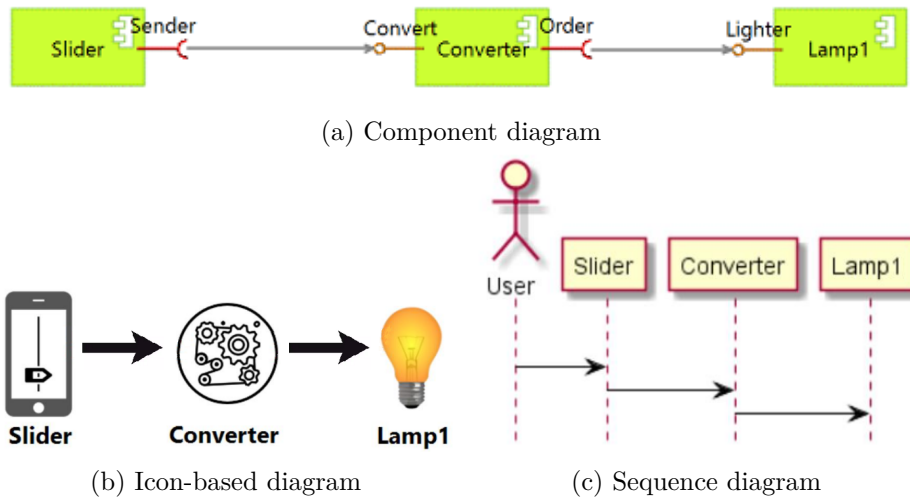


Figure 20: Possible views for an assembly in OCE

research questions. For each guideline, we discuss its significance and how it can enhance human interaction with an IML application.

In addition, as Amershi *et al.*'s framework is open for extension, we have incorporated two additional guidelines, GA and GB.

- *GA. Make clear how the system works.* This guideline emphasizes the importance of transparency regarding the AI's inner workings in order to gain trust from the human. While this may only be relevant for experienced users in complex applications, providing a simplified explanation of the inner workings can be beneficial for the interaction. As this guideline pertains to information that should be made available to the human at the beginning of the interaction, we have placed it in the "Initially" category.
- *GB. Provide interpretable output to the user.* This guideline emphasizes the need to ensure that the AI's output is understandable by the human. It is particularly relevant in complex use cases like Opportunistic Composition, whereas it may not be necessary in simpler scenarios such as content recommendation. This guideline pertains to information that is distilled during the interaction, and therefore we have put it in the "During Interaction" category.

Table 6 synthesizes the results of our study. There, S_{i-j} represents the proposed solution numbered j for the guideline G_i .

5.2.1 INITIALLY

These guidelines provide recommendations for the design of the beginning of the human-AI interaction, emphasizing the importance of managing the user's expectations and building trust in the intelligent service from the start.

G1. Make clear what the system can do. While the main purpose of an IML application, such as image classification, may be easy to understand, explaining the full range of capabilities to the user, such as data labeling or parameter tuning, can be more challenging.

This is particularly true in the case of an AmI application designed to assist humans in a pervasive environment like OCE. In this context, textual information (*S1-1*) or demonstrations (*S1-2*) may be the best options for explaining the application’s capabilities, *i.e.* building and proposing component assemblies to the human user. Over time and with use, the user should gain a better understanding of how the application works.

G2. Make clear how well the system can do what it can do .

The IML application should provide information to the human about the likelihood of approximations or mistakes occurring during the process. Managing human expectations is crucial to avoid over reliance on the AI system and prevent users from discontinuing use due to surprise or disappointment (Honeycutt et al., 2020).

In the case of reinforcement learning applications like OCE, mistakes are an essential part of the trial and error learning process, so providing clear and understandable warnings (*S2-1*) about this can initially help. Additionally, explaining what to do when the AI system produces incorrect output (*S2-2*), as discussed in Section 5.2.3, can further assist in managing expectations and maintaining user trust. Regarding OCE, it is crucial to communicate to the human that the assemblies proposed by OCE are built from the components currently available and the system’s past compositions.

GA. Make clear how the system works. This guideline emphasizes the importance of transparency in the AI system’s inner workings, which can help build trust with human users. Especially in complex applications, users want to know how the AI arrived at a particular decision or recommendation. Providing a simplified version of the inner workings can enhance the interaction, although it may only be necessary for experienced users.

In the context of OCE, there is currently no available explanation of how the engine works. Additionally, it may be difficult to demonstrate the improvement of the learner’s knowledge over time, as this is a gradual process that occurs over an extended period. To address this need for transparency, one potential solution is to provide information on the engine’s learning process. Textual explanations (*SA-1*) or targeted demonstrations (*SA-2*) should be considered.

5.2.2 DURING INTERACTION

These guidelines provide recommendations for designing the interaction itself. They primarily focus on managing the workload assigned to humans and addressing their social expectations.

G3. Time services based on context. Determining the appropriate moments to involve the human in the learning process is critical for a positive user experience. Disruptive interruptions can be detrimental to the user’s ability to complete tasks, leading them to disregard important information.

The OCE prototype currently generates a new assembly for every component that appears or disappears in the environment (*S3-1*), which aligns with this guideline. However, it may be beneficial to consider the user’s current focus and tailor the interactions accord-

ingly (*S3-2*), as suggested by (Kessler Faulkner & Thomaz, 2021). Nonetheless, detecting a user’s focus in an open AmI context can be more challenging than in a traditional desktop application.

G4. Show contextually relevant information. IML applications should display information relevant to the human’s current task or environment, to avoid overwhelming them with unnecessary information.

In the case of OCE, available components in the human’s environment are considered part of their context and are relevant to display (*S4-1*). However, taking into account additional information about the human’s context, such as their location or activity (*S4-2*), could further improve the relevance of the proposed assemblies. For example, if the human is in a classroom setting, it would be inappropriate to suggest applications for their favorite streaming service. However, automatically detecting the human’s context poses a significant challenge, especially for a generic application like OCE. Moreover, it should be noted that, fundamentally, OCE’s learning does not rely on learning patterns of behavior (Aztiria et al., 2010), activity detection or, more broadly, context identification.

GB. Provide interpretable output to the user. This recommendation is not included in the guidelines of Amershi et al. and the ability for users to intuitively understand the output of AI systems is not addressed by any other guideline. We believe this recommendation is crucial for AmI applications like OCE because such an understanding is essential for effective interaction with the AI system.

OCE’s first available view is a component diagram, which may not be suitable for end-users which are not trained in programming. In (Trouilhet et al., 2021), authors proposed to address this issue by using model transformation to provide alternative views that fit users’ comprehension capabilities (*SB-1*), as shown in Figure 20.

G5. Match relevant social norms. The interaction between the AI and human should be sensitive to the human’s social and cultural context to manage their expectations. For example, a voice assistant should address the human in a semi-formal manner, as it’s commonly expected in professional settings (Amershi et al., 2019).

However, in the case of OCE and ambient applications, it’s challenging to determine which social norms should be followed or avoided, as these applications may target a broad audience with diverse cultural backgrounds. At the current time, this guideline is not a priority to Opportunistic Composition.

G6. Mitigate social biases. The AI system should be free of any undesirable stereotypes or biases, which can arise during training and manifest in the model, particularly in applications related to race or gender (Gebru, 2019).

At present, OCE does not collect user data to identify common profiles, so there are no unfair biases of this nature that could emerge. Learning is only based on the available components in the human’s environment and the human’s actions within the application. While we acknowledge that biases and stereotypes can be present in any AI system, this particular guideline is not applicable to OCE as it stands. In the future, mitigating biases may become a concern if techniques such as transfer learning, as seen in Teachable Machine (Carney et al., 2020), are implemented.

5.2.3 WHEN WRONG

The following guidelines provide support for managing mistakes made by the learning service during human-AI interaction, when the output does not meet the expectations of the human. In most cases, the learning service tries to tune its service according to the preferences of the human, and as such, it needs to make mistakes to gain knowledge. Therefore, it is important to help the human recover from these mistakes to avoid disrupting their user experience.

G7. Support efficient invocation. The application should allow the human to easily request the AI's service. This means that if the output from the learner is not satisfactory, the human should be able to request new output quickly and easily.

In the current OCE workflow, a change in the environment triggers a new assembly to be pushed to the human, so the human doesn't explicitly request the AI's services. But the possibility for the human to trigger new assemblies at their discretion (*S7-1*) should be considered. Additionally, it would be possible for the human to express more precise needs in these composition requests, such as specifying components that should or should not be included in the final result (*S7-2*).

Alternatively, providing a blank editor would allow the human to freely edit component assemblies (*S7-3*). This would help them gain a better understanding of the way composite applications are built and provide an opportunity to gather data on the components they prefer or dislike, allowing to fine-tune the future propositions.

G8. Support efficient dismissal. The AI service should be easily dismissible or ignorable by the human, particularly when it produces incorrect results.

Currently, the OCE workflow allows the human to reject an assembly proposition from the composition engine (*S8-1*), which satisfies this guideline. However, further user testing is necessary to determine the effectiveness of this interaction and explore potential alternatives. In addition to rejection, we could also consider the option for the human to ignore OCE's assembly propositions (*S8-2*), which would be a lighter form of dismissal useful in situations where the human doesn't have time to review the proposed assembly. This action should have little to no impact on learning.

Another possible way for the human to dismiss OCE's proposition is to undo and rollback to a previous assembly (*S8-3*). Compared to basic rejection, this action would be a higher form of dismissal, and should send more negative feedback to OCE's agents. However, this may not always be possible in dynamic environments where components needed to deploy the previous assembly may no longer be available.

G9. Support efficient correction. Ensuring that the human can easily correct the AI when it is wrong is crucial for the success of the interaction.

In OCE, the correction process is already integrated, as the human can edit the links between components to modify an assembly (*S9-1*). However, this process requires a higher understanding of software composition and could be further improved to alleviate the workload on the human. For instance, we could consider highlighting alternative connections when the human wants to change the connection between two services (*S9-2*).

To explore alternative correction processes, we could look at existing solutions such as Crayons (Fails & Olsen Jr, 2003), which allows users to correct AI outputs by simply col-

oring over them (*S9-3*). Additional testing would be necessary to determine which solution is the most efficient and effective for the OCE workflow.

Guidelines G7, G8 and G9 propose to give the human the ability to choose between dismissing, ignoring, correcting, or reinvoking the AI service when it produces incorrect results. This flexibility allows the human to tailor their response to the severity of the mistake, which ultimately leads to a better user experience in the human-AI collaboration. However, it is important to consider the potential for placing an additional burden on the human when applying these guidelines.

G10. Scope services when in doubt. When uncertain about a human’s goals or preferences, the AI system should engage in disambiguation to ask for clarification. This not only acknowledges the human’s input, but also updates the system’s assumptions on the human’s preferences.

In the current OCE workflow, the system does not adapt to the level of uncertainty in the human’s preferences in the current environment. As a result, when the engine does not know what would be the best service for the human, it chooses at random.

To improve this, a recommendation is to adjust the system’s behavior based on the level of uncertainty. As an example, in situations where OCE is most uncertain and cannot differentiate between components, it could prompt the human to provide input by presenting them with a choice between the options (*S10-1*) (Flutura et al., 2018; Christiano et al., 2017). Alternatively, the system could propose multiple assemblies to explore different possibilities more quickly (*S10-2*). Ultimately, a minimal way to implement this guideline would be to at least inform the human of the uncertainty (*S10-3*), so that they can raise their levels of awareness towards any potential mistake from the engine.

To implement this, the uncertainty of the engine needs to be formalized, which can be challenging given OCE decentralized decision-making process. Managing a large number of large assemblies could also pose scaling problems for the human. Finally, this places an additional burden on the human user.

G11. Make clear why the system did what it did. Providing an explanation of AI’s behavior to the human is crucial in building trust and understanding of the AI application. However, creating explainable AIs is a current challenge for the scientific community (Arrieta et al., 2020).

The current OCE interface does provide access to the agent’s knowledge (*S11-1*), which can be considered an element of explanation for the agent’s decision-making process, as it is the state of its learning process. However, the current interface is not suitable for the general public and lacks information on the actual decision process, which is what the agent did with its knowledge. To meet this guideline, we need to work towards providing an easily understandable explanation of an agent’s decision to the human (*S11-2*). Additionally, by aggregating these individual explanations, we could create an emergent explanation of the application (*S11-3*) (Heuillet et al., 2022).

5.2.4 OVER TIME

These guidelines provide recommendations for how to improve the interaction over time, based on feedback and usage data. This can lead to increased user satisfaction, better

performance, and better trust in the system over time. Additionally, it can help to uncover potential issues or areas for improvement that may not have been apparent during the initial design and development phases.

G12. Remember recent interactions. The AI system should keep a memory of past interactions with the human to enhance efficiency.

OCE already takes into account the memory of past interactions and decisions made by the human, and gives a heavier weight to the most recent interaction (*S12-1*). Another way to see and implement this guideline would be to keep a memory of the last assemblies and allow the user to redeploy them if they are still available (*S12-2*). This could provide additional valuable feedback to the engine on the human’s preferences and help adapt to the changing environment.

G13. Learn from user behavior. The learner should personalize the human’s experience by learning from their actions, which is typically the main goal of an IML system.

In OCE, decisions are based on the past behavior of the human, such as accepting, rejecting, or modifying propositions of assemblies (*S13-1*). However, there are more sophisticated approaches proposed in the literature to better implement this guideline. For instance, in (Akrouf et al., 2014), authors suggest estimating the human’s error rate to improve recommendations (*S13-2*), while Zheng et al. (Zheng et al., 2018) consider the time to return to the service as a potential reward, where a faster return is considered better (*S13-3*). Implicit interaction data such as eye tracking, mouse movement, and timing of certain events can also be used to personalize the experience (*S13-4*). It is also possible to gather this information without informing the human, as in (Holzinger et al., 2019), to avoid overloading them with information.

G14. Update and adapt cautiously. An AI system should limit disruptive changes when updating and adapting the system’s behavior, in order not to shock and overload the human.

In OCE, this guideline is implemented through positive feedback given to assembly propositions that are accepted by the user, which encourages the engine to stick to those assembly plans over time and limit extreme changes in the propositions (*S14-1*). However, the dynamics of the environment may introduce changes in the available components, which is beyond the engine’s control. Moreover, in a reinforcement learning setting, some exploration is necessary, which can introduce changes in the system’s behavior. Nonetheless, since OCE is a multi-agent system, the exploration/exploitation dilemma is managed at the agent’s level, and the probability of all agents exploring simultaneously is very low, so that the proposed assemblies would not be overly disruptive. To further limit the risk of disruptive changes, it may be possible to work on alternative exploration strategies (*S14-2*), such as the one proposed in (Zheng et al., 2018).

G15. Encourage granular feedback. The human should be able to provide feedback at different levels of their interaction with the AI system.

The feedback in OCE is currently limited to the acceptance or rejection of proposed assemblies and editing of specific connections during modifications. To enrich the interaction, it would be interesting to allow the human to provide explicit feedback at all the other

levels of interaction shown in Figure 17, such as the assembly level (*S15-1*), component level (*S15-2*), connection level (*S15-3*), and service level (*S15-4*).

- At the assembly level, the human could provide explicit feedback on whether they like or dislike an assembly, which would be an alternative feedback to the simple acceptance or rejection of an assembly.
- At the component level, the human could provide feedback on which components they prefer to see in future assemblies.
- At the connection level, they could communicate their desire for certain connections, *i.e.*, component associations, to be prioritized or de-prioritized in future assemblies.
- At the service level, they could provide input on the services they want to see more or less utilized in future assemblies.

All those changes may induce changes in OCE’s behavior, since it is necessary to take into account those new kinds of feedback. All this range of feedback may allow the human to fine-tune the feedback they give to the engine, leading to better adaptation to their preferences and ultimately better propositions.

G16. Convey the consequences of user actions. It’s important for the human to be aware of how their actions can impact the future behavior of the AI system.

However, in the case of OCE, the human is not explicitly informed of this impact. As the proposed assemblies improve over time, the human may eventually notice how the system has evolved, but this may take some time to realize. Textual information like “you will see more/less similar applications in the future” could be an easy way to implement this guideline (*S16-1*). However, this may place a heavier cognitive load on the human, as they need to make choices based on their preferences (e.g. “do I want this application now”) and the impact of their choice on the system (e.g. “do I want to see more of this application in the future”).

Implementing more advanced propositions, such as the pre-visualization of future recommendations after a human action (*S16-2*), as suggested in (Schnabel et al., 2020), may be too difficult to achieve in a dynamic MAS setting where knowledge is distributed. In this context, it is also possible that some human users may not have the capacity to understand such complex information.

G17. Provide global controls. The human should be able to customize the overall functioning of the AI system based on their individual preferences.

The OCE system currently allows for this through editable engine parameters in the interface (*S17-1*). However, these parameters may not be understandable for end-users without ML expertise, making the system difficult to use. To address this, adjustments can be made to the interface, such as hiding the options to edit the engine’s global parameters initially (*S17-2*), to accommodate less experienced users (Carney et al., 2020). Furthermore, it is worth considering the concept of automatic parameter tuning (*S17-3*), which could reduce the workload on the human user.

G18. Notify user about changes. The AI system should inform the user when updates or new features are added, and provide guidance to existing users during the transition.

While this guideline does not currently apply to OCE since it is still a research prototype, it is an important consideration for future implementation. By keeping users informed about updates and changes to the system’s capabilities, users can feel more confident in their use of the system and understand how to take advantage of any new features or improvements.

5.3 Synthesis

In the previous section, we presented several design solutions taken from our analysis and from the literature presented in Section 3 that could help OCE adhere to the guidelines for human-AI interactions. These elements are summarized in Table 6.

As it stands, OCE does not comply with most of these guidelines, and several upgrades are needed to improve human-machine interaction.

In this section, we explore how an enhanced OCE that integrates the design solutions mentioned above could answer the research questions. Table 7 summarizes the level of response of this enhanced application, using the same scale of values as Table 4.

RQ11-Tasks. In this enhanced version of OCE, the human would still need to accept, reject or modify assemblies to give feedback to the engine. However, additional means of action, such as the optional ability to tune the global parameters (*S17-2*), or the ability to undo (*S8-3*) or ignore (*S8-2*) an assembly, would be possible. Since these additional tasks are optional, we rate the level of response of this version of OCE to RQ11 +, instead of ++ for the current version. This means that some ML-related tasks, like *S17-2*, are introduced into the interaction. Given that these tasks are optional, only experienced users with a strong understanding of the engine are expected to be concerned. Therefore, we believe this change in notation is acceptable.

RQ12-Workload. While most of the design solutions presented in the previous section allow the human to better understand the engine and how it works, they may require more attention from the human than a stripped-down user interface. For example, asking for disambiguation from the human when OCE is uncertain about how to proceed (*S10-1*) places an additional burden on the human. Although some ideas try to reduce this burden, such as not requesting too much input from the human (*S4-2*), it seems that the workload on the human would be quite heavy if OCE complied with all the guidelines. Therefore, we still rate it -. It is important to note that any additional workload in OCE is optional and dependent on the human’s skill level. Inexperienced users may choose to keep their interactions simple, while more experienced users can take on additional workloads as they become more familiar with the system.

RQ21-Information. Working with the guidelines enabled us to define several ways for the human to fine-tune the feedback they give to the learning mechanism, such as explicit feedback on components (*S15-2*), connections (*S15-3*), services (*S15-4*), or alternative ways of dismissing undesired output (*S8-2* and *S8-3*). With this additional information, the engine would better learn the preferences of the human, produce better assemblies, and ultimately provide a better user experience. Therefore, we rate it ++.

	Human-AI guidelines	Design solutions
Initially	G1. Make clear what the system can do.	<i>S1-1.</i> Textual explanations. <i>S1-2.</i> Demonstrations.
	G2. Make clear how well the system can do what it can do.	<i>S2-1.</i> Textual explanations. <i>S2-2.</i> Demonstrations on what to do when the engine is wrong.
	GA. Make clear how the system works.	<i>SA-1.</i> Textual explanations. <i>SA-2.</i> Demonstrations.
During interaction	G3. Time services based on context.	<i>S3-1.</i> New assembly at every change in the environment. <i>S3-2.</i> Request input when the human is attentive (Kessler Faulkner & Thomaz, 2021).
	G4. Show contextually relevant information.	<i>S4-1.</i> Display of the available components. <i>S4-2.</i> Automatically detected human activity.
	GB. Provide interpretable output to the user.	<i>SB-1.</i> Propose alternative views for an assembly (Trouilhet et al., 2021).
	G5. Match relevant social norms.	Not relevant for OCE.
	G6. Mitigate social biases.	Not relevant for OCE.
	G7. Support efficient invocation.	<i>S7-1.</i> Ability to request a new assembly. <i>S7-2.</i> New assembly with specifications. <i>S7-3.</i> Empty assembly editor.
When wrong	G8. Support efficient dismissal.	<i>S8-1.</i> Ability to reject an assembly. <i>S8-2.</i> Ability to ignore an assembly. <i>S8-3.</i> Ability to undo an assembly.
	G9. Support efficient correction.	<i>S9-1.</i> Ability to modify an assembly. <i>S9-2.</i> Highlights to assist correction. <i>S9-3.</i> An intuitive correction interface (Fails & Olsen Jr, 2003).
	G10. Scope services when in doubt.	<i>S10-1.</i> Human choice between components. <i>S10-2.</i> Multiple assemblies. <i>S10-3.</i> Notification of uncertainty.
	G11. Make clear why the system did what it did.	<i>S11-1.</i> Available agents' knowledge. <i>S11-2.</i> Agent's decision explanation. <i>S11-3.</i> Collective decision explanation.
	G12. Remember recent interactions.	<i>S12-1.</i> Give more weight to the most recent interaction. <i>S12-2.</i> Ability to redeploy the most recent assemblies.
Over time	G13. Learn from user behavior.	<i>S13-1.</i> Decisions based on past interactions. <i>S13-2.</i> Human error rate (Akrouf et al., 2014). <i>S13-3.</i> Complex interaction data (Zheng et al., 2018).
	G14. Update and adapt cautiously.	<i>S14-1.</i> Reinforcement learning limits disruptive changes. <i>S14-2.</i> Less disruptive exploration strategies (Zheng et al., 2018).
	G15. Encourage granular feedback.	<i>S15-1.</i> Assembly level feedback. <i>S15-2.</i> Component level feedback. <i>S15-3.</i> Connection level feedback. <i>S15-4.</i> Service level feedback.
	G16. Convey the consequences of user actions.	<i>S16-1.</i> Textual notifications. <i>S16-2.</i> Pre-visualization of an action's consequences (Schnabel et al., 2020).
	G17. Provide global controls.	<i>S17-1.</i> Editable engine parameters. <i>S17-2.</i> Optionally editable parameters (Carney et al., 2020). <i>S17-3.</i> Automatic parameter tuning.
	G18. Notify user about changes.	Not relevant for OCE.

Table 6: Synthesis of the proposed design solutions for each guideline for Human-AI interaction (Amershi et al., 2019). GA and GB were added for this analysis

RQ22-Assistance. Most of the guidelines tend to provide insights on how to better manage the human and provide assistance to them. In this enhanced version of OCE, the

Contribution	RQ11	RQ12	RQ21	RQ22
OCE with guidelines	+	-	++	++

Table 7: Summary of the levels of response to the different research questions for OCE after it applies the guidelines for human-AI interaction (the scale ”-” to ”++” indicates the level of response to a question)

human would be able to easily understand how the system works (*SA-1* and *SA-2*), be alerted when the engine is uncertain of its decisions (*S10-3*), and not be overloaded with unnecessary information (*S4-2*). Therefore, we rate the level of response to this question ++.

Scientific questions. We used the HAX toolkit to prioritize the design solutions presented above based on their feasibility and scientific value. While certain design solutions are solely engineering challenges that can be addressed and implemented, others pose scientific questions. Two of these questions are of particular interest: considering uncertainty (G10) and enhancing feedback granularity (G15).

The current OCE engine lacks the ability to estimate uncertainty. The incorporation of uncertainty estimation can enhance the reliability and trustworthiness of the system’s output to the human user. Addressing the challenge of integrating uncertainty into OCE requires not only accurate modeling but also the development of effective mechanisms for handling (*S10-1* or *S10-2*) and communicating (*S10-3*) uncertainty to the user.

Managing feedback granularity can lead to a more complete interaction between humans and the AI system. When the human can provide information at various levels of granularity, here spanning from the assembly (*S15-1*) to the service (*S15-4*), the learning mechanism receives more precise and detailed feedback on its performance, facilitating improved learning. Nevertheless, the integration of granular feedback is challenging and requires further research.

6. Conclusion

In this paper, we have explored the Interactive Machine Learning (IML) literature to identify potential solutions for the challenges faced by Ambient Intelligent (AmI) systems. We have conducted a critical review of the literature, resulting in a set of design solutions to improve human-AI interaction. Then we have considered their application in the case of Opportunistic Composition, and formulated 42 design solutions. In this concluding section, we discuss future directions and prospects for Opportunistic Composition, IML, and AmI based on the findings of this study.

6.1 Prospects for Opportunistic Composition

The next phase of our work is the development of a new prototype for OCE that implements the design solutions relative to the guidelines G10 and G15, which pose significant scientific challenges, as discussed in Section 5.3. This prototype will be put through a series of tests with human users to evaluate the impact of the design solutions on the user experience,

learning process, and performance. The results from these tests should provide important insights into the broader field of Interactive Machine Learning and Ambient Intelligence, and inform future research in these areas.

One important aspect to consider in the design of OCE is the diversity of human users' skills and knowledge. Some may have extensive technical expertise and be comfortable with programming and software development, while others may have limited experience in these areas. Therefore, it is essential to design OCE in a way that caters to the needs of all human users and allows them to participate at their own level of proficiency. Taking into account the diversity of users and their needs is an important scientific challenge in the development of OCE.

6.2 Prospects for Interactive Machine Learning

Simard et al. (Simard et al., 2017) make a parallel between Interactive Machine Learning and the general programming activity through their study on Interactive Machine Teaching (Section 3.2.1). IML and programming share, for example, the production by a human of an artifact (a learned model or a program) that meets the needs of one or more human users. This comparison leads the authors to believe that, just as programming has benefited from high-level tools and languages, like modern integrated development environments, IML needs its own tools and abstractions to facilitate the work of humans. Although the field of interactive machine learning and human-ML interaction is not yet mature, there are patterns and guidelines that can assist in this interaction. Our work may identify and contribute to these patterns.

6.3 Prospects for Ambient Intelligence

In light of the problems on Ambient Intelligence raised in Section 2.2, we looked to the guidelines for human-AI interaction (Amershi et al., 2019) as a framework to identify potential solutions. To ensure a smooth interaction between the human and the ambient system, we propose designing interaction modes that inform the human of the system's state and progress while taking into account their preferences. Providing customizable options and parameters will give more control to the human, enabling them to adjust the system to their liking.

Additionally, to dynamically adapt and evolve the intelligent behavior over time in response to changes in the user's goals and environment, the ambient system should learn from the human's actions and feedback. This allows the system to continually improve and adjust its behavior, while also taking into account any changes in the user's goals, preferences, and environment.

Lastly, to ensure relevant output given the inherent unpredictability of the environment, the system should be designed to handle uncertainty and ambiguity. One potential solution is employing algorithms that can reason under uncertainty and provide confidence estimates for their predictions. The human should provide feedback, particularly on the most uncertain answers, to enable the system to continually learn and improve.

6.4 Final Words

Although there is no one-size-fits-all solution for Human-AI interaction, we think that the findings from this study, particularly in the context of Ambient Intelligence, may be useful to the IML community and beyond. The insights gained from the literature review and analysis can inform the design of more effective human-AI interactions, as well as the development of new IML tools and abstractions to facilitate the work of humans. Furthermore, the implementation and testing of an OCE prototype incorporating the identified design solutions can provide empirical evidence of their effectiveness and further inform the development of human-AI systems. Ultimately, the goal is to create socially responsible (Cheng et al., 2021) intelligent systems that are not only capable of providing valuable assistance to humans, but also capable of doing so in a way that is intuitive, seamless, and tailored to the user’s individual needs and preferences. Achieving this goal will require continued collaboration and innovation across multiple fields, including AI, machine learning, human-computer interaction, and psychology, among others.

Data Availability

In this work, our analysis is founded upon the sources referenced within the text, in conjunction with the HAX toolkit, accessible at <https://www.microsoft.com/en-us/haxtoolkit/>. The source code for OCE, developed prior to this work, can be accessed at : <https://gitlab.irit.fr/oppocompo/oppocompo/OCE>.

Appendix A. Guidelines for Human-AI interaction

The following table presents the guidelines for Human-AI interaction by (Amershi et al., 2019). GA and GB were added for our analysis in Section 5.2.

References

- Aamir, M., & Bhusry, M. (2015). Recommendation system: state of the art approach. *International Journal of Computer Applications*, 120(12).
- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. (2016). TensorFlow: A System for Large-Scale Machine Learning. In *12th USENIX symposium on operating systems design and implementation (OSDI 16)*, pp. 265–283.
- Acampora, G., Cook, D. J., Rashidi, P., & Vasilakos, A. V. (2013). A survey on ambient intelligence in healthcare. *Proceedings of the IEEE*, 101(12), 2470–2494.
- Akrou, R., Schoenauer, M., Sebag, M., & Souplet, J.-C. (2014). Programming by feedback. In *Int. Conf. on Machine Learning*, Vol. 32, pp. 1503–1511. JMLR. org.
- Amershi, S., Cakmak, M., Knox, W. B., & Kulesza, T. (2014). Power to the people: The role of humans in interactive machine learning. *AI Magazine*, 35(4), 105–120.
- Amershi, S., Weld, D., Vorvoreanu, M., Fournery, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., & Inkpen, K. (2019). Guidelines for human-AI interaction. In *Proc. of the 2019 CHI conf. on human factors in computing systems*, pp. 1–13.

Guidelines for Human-AI interaction

Initially	<p>G1. Make clear what the system can do. Help the user understand what the AI system is capable of doing.</p> <p>G2. Make clear how well the system can do what it can do. Help the user understand how often the AI system may make mistakes.</p> <p>GA. Make clear how the system works. Help the user understand the inner workings of the AI system.</p>
During interaction	<p>G3. Time services based on context. Time when to act or interrupt based on the user’s current task and environment.</p> <p>G4. Show contextually relevant information. Display information relevant to the user’s current task and environment.</p> <p>GB. Provide interpretable output to the user. Help the user understand the AI system’s productions.</p> <p>G5. Match relevant social norms. Ensure the experience is delivered in a way that users would expect, given their social and cultural context.</p> <p>G6. Mitigate social biases. Ensure the AI system’s language and behaviors do not reinforce undesirable and unfair stereotypes and biases.</p>
When wrong	<p>G7. Support efficient invocation. Make it easy to invoke or request the AI system’s services when needed.</p> <p>G8. Support efficient dismissal. Make it easy to dismiss or ignore undesired AI system services.</p> <p>G9. Support efficient correction. Make it easy to edit, refine, or recover when the AI system is wrong.</p> <p>G10. Scope services when in doubt. Engage in disambiguation or gracefully degrade the AI system’s services when uncertain about a user’s goals.</p> <p>G11. Make clear why the system did what it did. Enable the user to access an explanation of why the AI system behaved as it did.</p>
Over time	<p>G12. Remember recent interactions. Maintain short term memory and allow the user to make efficient references to that memory.</p> <p>G13. Learn from user behavior. Personalize the user’s experience by learning from their actions over time.</p> <p>G14. Update and adapt cautiously. Limit disruptive changes when updating and adapting the AI system’s behaviors.</p> <p>G15. Encourage granular feedback. Enable the user to provide feedback indicating their preferences during regular interaction with the AI system.</p> <p>G16. Convey the consequences of user actions. Immediately update or convey how user actions will impact future behaviors of the AI system.</p> <p>G17. Provide global controls. Allow the user to globally customize what the AI system monitors and how it behaves.</p> <p>G18. Notify user about changes. Inform the user when the AI system adds or updates its capabilities.</p>

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., et al. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115.

Aztiria, A., Izaguirre, A., & Augusto, J. C. (2010). Learning patterns in ambient intelligence environments: a survey. *Artificial Intelligence Review*, 34, 35–51.

Berg, S., Kutra, D., Kroeger, T., Straehle, C. N., Kausler, B. X., Haubold, C., Schiegg, M., Ales, J., Beier, T., & Rudy, M. (2019). Ilastik: interactive machine learning for (bio) image analysis. *Nature Methods*, 16(12), 1226–1232.

- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., & Zaremba, W. (2016). OpenAI Gym. arXiv preprint arXiv:1606.01540.
- Carney, M., Webster, B., Alvarado, I., Phillips, K., Howell, N., Griffith, J., Jongejan, J., Pitaru, A., & Chen, A. (2020). Teachable machine: Approachable Web-based tool for exploring machine learning classification. In *Extended abstracts of the 2020 CHI Conf. on human factors in computing systems*, pp. 1–8.
- Celemin, C., & Ruiz-del Solar, J. (2019). An interactive framework for learning continuous actions policies based on corrective feedback. *Journal of Intelligent & Robotic Systems*, 95(1), 77–97.
- Chatzimparmpas, A., Martins, R. M., Jusufi, I., & Kerren, A. (2020). A survey of surveys on the use of visualization for interpreting machine learning models. *Information Visualization*, 19(3), 207–233.
- Cheng, L., Varshney, K. R., & Liu, H. (2021). Socially responsible AI algorithms: Issues, purposes, and challenges. *Journal of Artificial Intelligence Research*, 71, 1137–1181.
- Christiano, P., Leike, J., Brown, T. B., Martic, M., Legg, S., & Amodei, D. (2017). Deep reinforcement learning from human preferences..
- Delcourt, K., Adreit, F., Arcangeli, J.-P., Hacid, K., Trouilhet, S., & Younes, W. (2021). Automatic and Intelligent Composition of Pervasive Applications - Demonstration. In *19th IEEE Int. Conf. on Pervasive Computing and Communications (PerCom 2021)*, Kassel (virtual), Germany.
- Dellermann, D., Calma, A., Lipusch, N., Weber, T., Weigel, S., & Ebel, P. (2019). The Future of Human-AI Collaboration: A Taxonomy of Design Knowledge for Hybrid Intelligence Systems. In *Proceedings of the 52nd Hawaii Int. Conf. on System Sciences*.
- Dorigo, M., Birattari, M., & Stutzle, T. (2006). Ant colony optimization. *IEEE Computational Intelligence magazine*, 1(4), 28–39.
- Dudley, J. J., & Kristensson, P. O. (2018). A review of user interface design for interactive machine learning. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 8(2), 1–37.
- Dunne, R., Morris, T., & Harper, S. (2021). A survey of ambient intelligence. *ACM Computing Surveys (CSUR)*, 54(4), 1–27.
- Fails, J. A., & Olsen Jr, D. R. (2003). Interactive machine learning. In *Proceedings of the 8th Int. Conf. on Intelligent user interfaces*, pp. 39–45.
- Flutura, S., Seiderer, A., Aslan, I., Dang, C.-T., Schwarz, R., Schiller, D., & André, E. (2018). Drinkwatch: A mobile wellbeing application based on interactive and cooperative machine learning. In *Proceedings of the 2018 Int. Conf. on Digital Health*, pp. 65–74.
- Geburu, T. (2019). Oxford handbook on AI ethics, book chapter on race and gender. arXiv preprint arXiv:1908.06165.
- Geva, S., & Sitte, J. (1993). A cartpole experiment benchmark for trainable controllers. *IEEE Control Systems Magazine*, 13(5), 40–51.

- Grant, M. J., & Booth, A. (2009). A typology of reviews: an analysis of 14 review types and associated methodologies. *Health information & libraries journal*, 26(2), 91–108.
- Heuillet, A., Couthouis, F., & Díaz-Rodríguez, N. (2022). Collective explainable AI: Explaining cooperative strategies and agent contribution in multiagent reinforcement learning with shapley values. *IEEE Computational Intelligence Magazine*, 17(1), 59–71.
- Holzinger, A., Plass, M., Kickmeier-Rust, M., Holzinger, K., Crişan, G. C., Pintea, C.-M., & Palade, V. (2019). Interactive machine learning: experimental evidence for the human in the algorithmic loop. *Applied Intelligence*, 49(7), 2401–2414.
- Honeycutt, D., Nourani, M., & Ragan, E. (2020). Soliciting human-in-the-loop user feedback for interactive machine learning reduces user trust and impressions of model accuracy. In *Proceedings of the AAAI Conf. on Human Computation and Crowdsourcing*, Vol. 8, pp. 63–72.
- Hutter, F., Kotthoff, L., & Vanschoren, J. (2019). *Automated machine learning: methods, systems, challenges*. Springer Nature.
- Jiang, L., Liu, S., & Chen, C. (2019). Recent research advances on interactive machine learning. *Journal of Visualization*, 22(2), 401–417.
- Karmaker, S. K., Hassan, M. M., Smith, M. J., Xu, L., Zhai, C., & Veeramachaneni, K. (2021). Automl to date and beyond: Challenges and opportunities. *ACM Computing Surveys (CSUR)*, 54(8), 1–36.
- Kessler Faulkner, T. A., & Thomaz, A. (2021). Interactive Reinforcement Learning from Imperfect Teachers. In *Companion of the 2021 ACM/IEEE Int. Conf. on Human-Robot Interaction*, pp. 577–579.
- Li, G., Gomez, R., Nakamura, K., & He, B. (2019). Human-centered reinforcement learning: A survey. *IEEE Transactions on Human-Machine Systems*, 49(4), 337–349.
- Liang, C., Proft, J., Andersen, E., & Knepper, R. A. (2019). Implicit communication of actionable information in human-AI teams. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–13.
- Makonin, S., Bartram, L., & Popowich, F. (2012). A smarter smart home: Case studies of ambient intelligence. *IEEE pervasive computing*, 12(1), 58–66.
- O’Neill, T., McNeese, N., Barron, A., & Schelble, B. (2022). Human–autonomy teaming: A review and analysis of the empirical literature. *Human Factors*, 64(5), 904–938.
- Quinlan, J. R. (1986). Induction of decision trees. *Machine learning*, 1(1), 81–106.
- Ramos, G., Meek, C., Simard, P., Suh, J., & Ghorashi, S. (2020). Interactive machine teaching: a human-centered approach to building machine-learned models. *Human-Computer Interaction*, 35(5-6), 413–451.
- Robinson, D. C., Sanders, D. A., & Mazharsolook, E. (2015). Ambient intelligence for optimal manufacturing and energy efficiency. *Assembly Automation*, 35(3), 234–248.
- Sadri, F. (2011). Ambient intelligence: A survey. *ACM Computing Surveys*, 43(4), 1–66.

- Saisubramanian, S., Kamar, E., & Zilberstein, S. (2022). Avoiding negative side effects of autonomous systems in the open world. *Journal of Artificial Intelligence Research*, *74*, 143–177.
- Schnabel, T., Amershi, S., Bennett, P. N., Bailey, P., & Joachims, T. (2020). The Impact of More Transparent Interfaces on Behavior in Personalized Recommendation. In *Proceedings of the 43rd Int. ACM SIGIR Conf. on Research and Development in Information Retrieval*, pp. 991–1000.
- Simard, P. Y., Amershi, S., Chickering, D. M., Pelton, A. E., Ghorashi, S., Meek, C., Ramos, G., Suh, J., Verwey, J., & Wang, M. (2017). Machine teaching: A new paradigm for building machine learning systems. arXiv preprint arXiv:1707.06742.
- Stankovic, J. A., Ma, M., Preum, S. M., & Alemzadeh, H. (2021). Challenges and directions for ambient intelligence: A cyber physical systems perspective. In *2021 IEEE Third International Conference on Cognitive Machine Intelligence (CogMI)*, pp. 232–241. IEEE.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Trouilhet, S., Arcangeli, J.-P., Bruel, J.-M., & Koussaifi, M. (2021). Model-driven engineering for end-users in the loop in smart ambient systems. *Journal of Universal Computer Science*, *27*(7), 755–773.
- Van den Bos, K., Vermunt, R., & Wilke, H. A. (1996). The consistency rule and the voice effect: The influence of expectations on procedural fairness judgements and performance. *European Journal of Social Psychology*, *26*(3), 411–428.
- Velastin, S. A., Boghossian, B. A., Lo, B. P., Sun, J., & Vicencio-Silva, M. A. (2004). Prismatic: toward ambient intelligence in public transport environments. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, *35*(1), 164–182.
- Wall, E., Ghorashi, S., & Ramos, G. (2019). Using expert patterns in assisted interactive machine learning: A study in machine teaching. In *IFIP Conf. on Human-Computer Interaction*, pp. 578–599. Springer.
- Weiss, K., Khoshgoftaar, T. M., & Wang, D. (2016). A survey of transfer learning. *Journal of Big data*, *3*(1), 1–40.
- Yin, M., Wortman Vaughan, J., & Wallach, H. (2019). Understanding the effect of accuracy on trust in machine learning models. In *Proceedings of the 2019 chi conference on human factors in computing systems*, pp. 1–12.
- Younes, W., Trouilhet, S., Adreit, F., & Arcangeli, J.-P. (2020). Agent-mediated application emergence through reinforcement learning from user feedback. In *29th IEEE International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises*.
- Zhang, R., McNeese, N. J., Freeman, G., & Musick, G. (2021a). "An Ideal Human": Expectations of AI Teammates in Human-AI Teaming. *Proceedings of the ACM on Human-Computer Interaction*, *4*(CSCW3), 1–25.

- Zhang, R., Torabi, F., Warnell, G., & Stone, P. (2021b). Recent advances in leveraging human guidance for sequential decision-making tasks. *Autonomous Agents and Multi-Agent Systems*, 35(2), 1–39.
- Zheng, G., Zhang, F., Zheng, Z., Xiang, Y., Yuan, N. J., Xie, X., & Li, Z. (2018). Drn: A deep reinforcement learning framework for news recommendation. In *Proceedings of the 2018 world wide web conference*, pp. 167–176.