



<http://www.diva-portal.org>

## Postprint

This is the accepted version of a chapter published in *Engineering multi-agent systems: 6th international workshop, EMAS 2018, Stockholm, Sweden, July 14-15, 2018, revised selected papers*.

Citation for the original published chapter:

Kampik, T., Nieves, J C., Lindgren, H. (2019)

Empathic autonomous agents

In: Danny Weyns, Viviana Mascardi and Alessandro Ricci (ed.), *Engineering multi-agent systems: 6th international workshop, EMAS 2018, Stockholm, Sweden, July 14-15, 2018, revised selected papers* (pp. 181-201). Cham: Springer

Lecture Notes in Artificial Intelligence

[https://doi.org/10.1007/978-3-030-25693-7\\_10](https://doi.org/10.1007/978-3-030-25693-7_10)

N.B. When citing this work, cite the original published chapter.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:umu:diva-161633>

# Empathic Autonomous Agents

Timotheus Kampik<sup>[0000-0002-6458-2252]</sup>, Juan Carlos Nieves<sup>[0000-0003-4072-8795]</sup>,  
and Helena Lindgren<sup>[0000-0002-8430-4241]</sup>

Umeå University, 901 87, Umeå, Sweden  
{tkampik, jcnieves, helena}@cs.umu.se

**Abstract.** Identifying and resolving conflicts of interests is a key challenge when designing autonomous agents. For example, such conflicts often occur when complex information systems interact persuasively with humans and are in the future likely to arise in non-human agent-to-agent interaction. We introduce a theoretical framework for an *empathic* autonomous agent that proactively identifies potential conflicts of interests in interactions with other agents (and humans) by considering their utility functions and comparing them with its own preferences using a system of shared values to find a solution all agents consider *acceptable*. To illustrate how empathic autonomous agents work, we provide running examples and a simple prototype implementation in a general-purpose programming language. To give a high-level overview of our work, we propose a reasoning-loop architecture for our empathic agent.

**Keywords:** Multi-agent systems · Utility theory · Conflicts of interests

## 1 Background and Problem Description

In modern information technologies, conflicts of interests between users and information systems that operate with a high degree of autonomy (*autonomous agents*) are of increasing prevalence. For example, complex web applications persuade end-users, possibly against the interests of the persuaded individuals<sup>1</sup>. Given the prevalence of autonomous systems will increase, conflicts between autonomous agents and humans (or between different autonomous agent instances and types) can be expected to occur more frequently in the future, e.g. in interactions with or among autonomous vehicles in scenarios that cannot be completely solved by applying static traffic rules. Consequently, one can argue for the need to develop *empathic* intelligent agents that consider the preferences or utility functions of others, as well as ethics rules and social norms when interacting with their environment to avoid severe conflicts of interests. As a simple example, take two vehicles (*A* and *B*) that are about to enter a bottleneck. Assume they cannot enter the bottleneck at the same time. *A* and *B* can either wait or drive. Considering only its own utility function, *A* might determine that *driving* is the best action to execute, given that *B* will likely stop and wait to avoid a crash. However, *A* should ideally assess both its own and *B*'s utility function and act accordingly. If *B*'s

<sup>1</sup> E.g., research provides evidence that contextual advertisement influences how users process online news [25]; social network applications have effectively been employed for political persuasion (see for an example: [4]).

utility for driving is considered higher than A's, A can then come to the conclusion that *waiting* is the best action. As A does not only consider its own goals, but also the ones of B, one can regard A as *empathic*, following Coplan's definition of empathy, as "a process through which an observer simulates another's situated psychological states, while maintaining clear self–other differentiation" [12]. While existing literature covers conflict resolution in multi-agent systems from a broad range of perspectives (see for a partial overview: [2]), devising a theoretical framework for autonomous agents that consider the utility functions (or preferences) of agents in their environment and use a combined utilitarian/rule-based approach to identify and resolve conflicts of interests can be considered a novel idea. However, existing multi-agent systems research can be leveraged to implement core components of such a framework, as is discussed later.

In this chapter, we provide the following research contributions:

1. We create a theoretical framework for an empathic agent that uses a combination of utility-based and rule-based concepts to compromise with other agents in its environment when deciding upon how to act.
2. We provide a set of running examples that illustrate how the empathic agent works and show how the examples can be implemented in a general-purpose programming language.
3. We propose a reasoning-loop architecture for a generic empathic agent.

The rest of this chapter is organized as follows: in Section 2, we present a theoretical framework for the problem in focus. Then, we illustrate the concepts with the help of different running examples and describe the example implementation in a general-purpose programming language in Section 3. Next, we outline a basic reasoning-loop architecture for the empathic agent in Section 4. In Section 5, we analyze how the architecture aligns with the belief-desire-intention approach and propose an implementation using the Jason multi-agent development framework. Finally, we discuss how our empathic agent concepts relate to existing work, propose potential use cases, highlight a set of limitations, and outline future work in Section 6, before we conclude the chapter in Section 7.

## 2 Empathic Agent Core Concepts

In this section, we describe the core concepts of the empathic agent. To allow for a precise description, we assume the following scenario<sup>2</sup>:

- The scenario describes the interaction between a set of empathic agents  $\{A_0, \dots, A_n\}$ .
- Each interaction scenario takes place at one specific point in time, at which all agents execute their actions simultaneously.
- At this point in time, each agent  $A_i (0 \leq i \leq n)$  has a finite set of *possible* actions  $Acts_i := \{Act_i^0, \dots, Act_i^m\}$ , resulting in an overall set of action sets  $Acts := \{Acts_0, \dots, Acts_n\}$ . Each agent can execute an *action tuple* that contains one or

<sup>2</sup> As we will explain later, the scenario and the resulting specification can be gradually extended to allow for better real-world applicability.

multiple actions. In each interaction scenario, all agents execute their actions simultaneously and receive their utility as a numeric reward based on the actions that have been executed.

- The utility of an agent  $A_i$  is determined by a function  $u_i$  of the actions of all agents. The utility function returns a numerical value or *null*<sup>3</sup>:

$$u_i := Acts_0 \times \dots \times Acts_n \rightarrow \{null, -\infty, \mathbb{R}, \infty\}$$

The goal of the empathic agent is to maximize its own utility as long as no conflicts with other agents arise. We define a conflict of interests between several agents as any interaction scenario in which there is no tuple of possible actions that maximizes the utility functions of all agents. I.e., we need to compare  $\arg \max u_{A_0}, \dots, \arg \max u_{A_n}$ <sup>4</sup>. Note that  $\arg \max u_{A_i}$  returns a *set of tuples* (that contains all action tuples that yield the maximal utility for agent  $A_i$ ). For this, we create a boolean function  $c$  that the empathic agent uses to determine conflicts between itself and other agents, based on the utility functions of all agents:

$$c(u_{A_0}, \dots, u_{A_n}) := \begin{cases} true, & \text{if } \\ & \arg \max u_{A_0} \cap \dots \cap \arg \max u_{A_n} \neq \{\}; \\ false, & \text{otherwise.} \end{cases}$$

Considering the incomparability property of the von Neumann-Morgenstern utility theorem [24], such a conflict can be solved only if a system of values exists that is shared between the agents and used to determine *comparable* individual utility values. Hence, we introduce such a shared value system. To provide a possible structure for this system, we deconstruct the utility functions into two parts:

- An actions-to-consequences mapping (a function  $a2c_i$  that takes the actions the agents potentially decide to execute and returns a set of *consequences* (propositional atoms)  $Consqs := \{Consq_i^0, \dots, Consq_i^n\}$ ):

$$a2c_i := Acts_i \times \dots \times Acts_n \rightarrow 2^{Consqs}$$

- A consequences-to-utility mapping (utility quantification function  $uq$ ). Note that the actions-to-consequences mapping is agent-specific, while the utility quantification function is generically provided by the shared value system<sup>5</sup>:

$$uq := 2^{Consqs} \rightarrow \{null, -\infty, \mathbb{R}, \infty\}$$

<sup>3</sup> We allow for utility functions to return a *null* value for action tuples that are considered impossible, e.g. in case some actions are mutually exclusive. While we concede that the elegance of this approach is up for debate, we opted for it because of its simplicity.

<sup>4</sup> The  $\arg \max$  operator takes the function it precedes and returns all argument tuples that maximize the function.

<sup>5</sup> I.e., for the same actions, an agent should only receive a different utility outcome than another agent if the impact on the two is distinguishable in its *consequences*. We again allow for *null* values to be returned in case of impossible action tuples.

Then, agents can agree on the utility value of a given tuple of actions, as long as the *quality* of the consequence is observable to all agents in the same way. In addition, the value system can introduce generally applicable rules, e.g. to hard-code a prioritization of individual freedom into an agent. With help of the value system, we create a pragmatic definition of a conflict of interests as any situation, in which there is no tuple of actions that is regarded as *acceptable* by all agents when considering the shared set of values, given each agent executes the actions that maximize their individual utility function. To support the notion of *acceptability*, we introduce a set of agent-specific *acceptability functions*  $accs := \{acc_{A_0}, \dots, acc_{A_n}\}$ . The acceptability functions are derived from the corresponding utility functions and the shared system of values and take a set of actions as their inputs. Acceptability functions are domain-specific and there is no generic logic to be described in this context:

$$acc_{A_i} := Acts_{A_0} \times \dots \times Acts_{A_n} \rightarrow \{null, true, false\}$$

The notion of acceptability rules adds a *normative* aspect to the otherwise *consequentialist* empathic agent framework. Without this notion, our definition of a conflict of interests would cover many scenarios that most human societies regard as not conflict-worthy, e.g. when one agent would need to accept large utility losses to optimize its own actions towards improving another agents' utility. Considering the acceptability functions, we can now determine whether a conflict of interests in terms of the *pragmatic* definition approach exists for an agent  $A_i$  by using the following function  $cp$  that takes the utility function  $u_i$  of agent  $A_i$  and the acceptability functions  $Accs := \{acc_{A_0}, \dots, acc_{A_n}\}$  as input arguments:

$$cp(u_i, Accs) := \begin{cases} true, & \text{if:} \\ \quad \nexists acts \in \arg \max u_i \wedge \forall acc \in Accs : acc(acts) = true \\ false, & \text{otherwise.} \end{cases}$$

We define an *empathic* agent  $A_i$  as an agent that, when determining the actions it executes, considers the utility functions of the agents it could potentially affect and maximizes its own utility only if doing so does not violate the acceptability function of any other agent; otherwise it acts to maximize the shared utility of all agents (while also considering the acceptability functions)<sup>6</sup>. Algorithm 1 specifies an initial, *naive* approach towards the empathic agent core algorithm. The empathic agent core algorithm of an agent  $A_i$  in its simplest form can be defined as a function that takes the utility functions  $\{u_0, \dots, u_n\}$  of the different agents, the set of all acceptability functions  $Accs := \{acc_0, \dots, acc_1\}$ , and all possible actions  $Acts_i$  of agent  $A_i$  and returns the tuple of actions  $A_i$  should execute<sup>7</sup>.

<sup>6</sup> As different aggregation approaches are possible (for example: *sum*, *product*) to determine the maximal shared utility, we introduce the not further specified aggregation function  $aggregate(u_0, \dots, u_n)$ . In our running examples (see Section 3), we use the product of the individual utility function outcomes to introduce some notion of fairness; inequality should not be in the interest of the empathic agent. However, the design choice for this implementation detail can be discussed.

<sup>7</sup> To facilitate readability, we switch to a pseudo-code notation for the following algorithms.

---

**Algorithm 1** Naive empathic agent algorithm: *D\_A\_N* (*determine actions naive*)

---

```

1: procedure D_A_Ni({u0, ..., un}, Accs, Actsi)    ▷ Utility & acceptability functions of all
   agents, actions of Ai (0 ≤ i ≤ n)
2:   if ∃acts ∈ arg max ui ∧ ∀acc ∈ Accs : acc(acts) = true then
3:     best_acceptable_acts ← ⋃acts ∈ arg max ui : ∀acc ∈ Accs acc(acts) = true
4:     return Actsi ∩ first(actsk ∈ best_acceptable_acts)
5:   else
6:     return Actsi ∩ first(arg max(aggregate(u'0, ..., u'n)))
7:   end if
8: end procedure

```

---

Note that in the context of the empathic agent algorithms, the function *first*(*set*) turns the provided *set of tuples* into a *sequence of tuples* by sorting the elements in decreasing alphanumerical order and then returns the first element of the sequence. This enables a deterministic action tuple selection. Moreover, we construct a set of new utility functions {*u*'<sub>0</sub>, ..., *u*'<sub>*n*</sub>} that assign all not acceptable action tuples a utility of *null* (Algorithm 2)<sup>8</sup>:

---

**Algorithm 2** Helper function: new utility function based on *u*<sub>*i*</sub>; all not acceptable action tuples yield utility of *null*.

---

```

1: procedure u'i(ui, {acts0, ..., actsn}, accs)
2:   is_acceptable ← ∀acc ∈ accs : acc(actsi) = true
3:   if is_acceptable then
4:     return ui(actsi, ..., actsn)
5:   else
6:     return null
7:   end if
8: end procedure

```

---

In Algorithm 1, we specify that the agent picks the first item in the sequence of determined action tuples if it finds multiple optimal tuples of actions. Alternatively, the agent could employ one of the following approaches to select between the optimal action tuples:

- **Random.** The agent picks a random action tuple from the list of the tuples it determined as optimal. This would require empathic agents to use an additional protocol to agree on the action tuple that should be executed.
- **Utilitarian.** Among the action tuples that were determined as optimal, the agent picks the one that provides maximal combined utility for all agents and falls back to a random or first-in-sequence selection between action tuples if several of such tuples exist.

---

<sup>8</sup> We already use *null* to denote *impossible* action tuples. This implies an acceptable action tuple should always exist. To achieve a distinction, a value of  $-\infty$  could be assigned.

Still, the algorithm is somewhat naive, as agents that implement it will decide to execute suboptimal activities if the following conditions apply:

- Multiple agents find that the actions that optimize their individual utility are inconsistent with the actions that are optimal for at least one of the other agents.
- Multiple agents find that executing these *conflicting* actions is considered *acceptable*.
- Executing these *acceptable* actions generates a lower utility for both agents than optimizing the shared utility would.

Hence, we extend the algorithm so that the agent selects the tuple of actions that maximizes its own utility, but falls back to maximize shared utility if the utility-maximizing action tuple is either not acceptable, or would lead to a lower utility outcome than maximizing the shared utility, considering the other agent follows the same approach (Algorithm 3):

---

**Algorithm 3** Lazy empathic agent algorithm: *D\_AL* (*determine actions lazy*)

---

```

1: procedure D_AL( $\{u_0, \dots, u_n\}, Accs$ )  $\triangleright$  Utility & acceptability functions of all agents,
   actions of all agents  $\{A_0, \dots, A_n\}$ 
2:    $\{acts\_max_0, \dots, acts\_max_n\} \leftarrow DETERMINE\_ACT\_MAX(u_i, Accs)$ 
3:    $\{good\_acts\_max_0, \dots, good\_acts\_max_n\} \leftarrow \{$ 
4:      $DETERMINE\_GOOD\_ACTS\_MAX(u_0, Accs, acts\_max_0),$ 
5:      $\dots,$ 
6:      $DETERMINE\_GOOD\_ACTS\_MAX(u_n, Accs, acts\_max_n),$ 
7:    $\}$ 
8:   if  $good\_acts\_max_0 \cap \dots \cap good\_acts\_max_n \neq \{\}$  then
9:     return  $Acts_i \cap first(good\_acts\_max)$ 
10:  else
11:    return  $Acts_i \cap first(\arg \max(aggregate(u'_0, \dots, u'_n)))$ 
12:  end if
13: end procedure

```

---

Algorithm 3 calls two helper functions. Algorithm 4 determines *acceptable* action tuples that maximize a provided utility function  $u_i$ :

---

**Algorithm 4** Helper function: determine acceptable action tuples that maximize utility function  $u_i$

---

```

1: procedure DETERMINE_ACT_MAX( $u_i, Accs$ )
2:   return  $\bigcup_{acts \in \arg \max u_i} : \forall acc \in Accs acc(acts) = true$ 
3: end procedure

```

---

Algorithm 5 determines all action tuples that would maximize an agent's ( $A_i$ 's) utility if this agent could *dictate* the actions of all other agents, given the action tuples provide

a better utility for this agent than the action tuples that maximize all agents' combined utility, given all agents execute an action tuple that maximizes their own utility if they could dictate the other agents' actions. Note that Algorithm 5 makes use of the previously introduced algorithm (Algorithm 1):

---

**Algorithm 5** Helper function: determines all maximizing action tuples that would still yield a good utility result for agent  $A_i$  ( $0 \leq i \leq n$ ), given all other agents also pick an action tuple that would maximize their own utility, if all other agents “played along”.

---

```

1: procedure DETERMINE_GOOD_ACTS_MAX( $u_i, Accs$ )
2:   return  $\bigcup_{acts \in \arg \max u_i} : \forall acc \in Accs : acc(acts) = true \wedge$ 
3:          $u_i(\bigcup_{k=0}^n D\_A\_N_k(\{u_0, \dots, u_n\}, Accs, acts))$ 
4:          $\geq u_i(acts\_max)$ 
5: end procedure

```

---

However, this algorithm only considers two types of action tuples for execution: action tuples that provide the maximal individual utility for the agent and action tuples that provide the maximal combined utility for all agents. Action tuples that do not maximize the agent's individual utility, but are still preferable over the action tuples that maximize the combined utility, remain unconsidered. Consequently, we call an agent that implements such an algorithm a *lazy* empathic agent. We extend the algorithm to also consider all action tuples that could possibly be relevant. I.e., if an action tuple is not considered acceptable, or if the tuple is considered acceptable but the agent chooses to not execute it, the agent falls back to the tuple of actions that provides the next best individual utility. We construct a function *ne* that returns the *Nash equilibria* based on the updated utility functions  $\{u'_0, \dots, u'_n\}$ , considering we have a strategic game  $\langle N, (A_i) \succsim_i \rangle$ , with  $N := \{A_0, \dots, A_n\}$ ,  $A_i := Acts_{A_i}$ , and  $acts \succsim_i acts' := u'_i(acts) \geq u'_i(acts')$ <sup>9</sup>. Then, we create the *full* empathic agent core algorithm  $D\_A\_F_i$  for an agent  $A_i$  that takes the updated utility functions  $\{u'_0, \dots, u'_n\}$  and all agents' possible actions as inputs  $\{Acts_0, \dots, Acts_n\}$ . The algorithm determines the (first of) the Nash equilibria that provide the highest shared utility and, if no Nash equilibrium exists, chooses the first tuple of actions that maximizes shared utility:

---

<sup>9</sup> See the Nash equilibrium definition provided by Osborne and Rubinstein [19, p. 11 et sqq.].

**Algorithm 6** Full empathic agent algorithm:  $D\_A\_F$  (*determine actions full*)

---

```

1: procedure  $D\_A\_F_i(\{u'_0, \dots, u'_n\}, \{Acts_0, \dots, Acts_n\})$ 
2:    $equilibria \leftarrow ne(\{u'_0, \dots, u'_n\}, \{Acts_0, \dots, Acts_n\})$ 
3:   if  $equilibria \neq \{\}$  then
4:      $shared\_max\_equilibria \leftarrow acts^* \in equilibria :$ 
5:        $\forall acts \in equilibria :$ 
6:          $(u'_0(acts^*) \times \dots \times u'_n(acts^*)) \geq (u'_0(acts) \times \dots \times u'_n(acts))$ 
7:       return  $Acts_i \cap first(shared\_max\_equilibria)$ 
8:   else
9:     return  $Acts_i \cap first(\arg \max(aggregate(u_0, \dots, u_n)))$ 
10:
11:   end if
12: end procedure

```

---

Going back to the selection between several action tuples that might be determined as optimal, it is now clear that a *deterministic* approach for selecting a final action tuple is preferable for both *lazy* and *full* empathic agents, as it avoids agents deciding upon executing action tuples that are not aligned with one another and lead to an unnecessary low utility outcome. Hence, we propose using a *utilitarian* approach with a first-in-sequence selection if the utilitarian approach is inconclusive<sup>10</sup>.

The proposed agent can be considered a *rational agent* following the definition by Russel and Norvig in that it “acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome” [22, p. 4-5] and an *artificially socially intelligent agent* as defined by Dautenhahn as it instantiates “human-style social intelligence” in that it “manage[s] the individual’s [its own] interests in relationship to the interests of the social system of the next higher level” [13].

### 3 Running Examples

In this section, we present two simple running examples of empathic agents and describe the implementation of the examples in a general-purpose programming language (JavaScript).

#### 3.1 Example 1: Vehicles

We provide a running example for the “vehicle/bottleneck” scenario introduced above. Consequently, we have a two-agent scenario  $\{A, B\}$ . Each agent has a utility function  $u_{A,B} := Acts_A \times Acts_B \rightarrow \{-\infty, \mathbb{R}, \infty\}$ .  $Acts_A$  and  $Acts_B$  are the possible actions A and B, respectively, can execute. To fully specify the utility functions, we follow the approach outlined above and first construct the actions-to-consequences mappings  $a2c_A$  and  $a2c_B$  for both agents. The possible actions are  $Acts_A = \{drive_A, wait_A\}$  and  $Acts_B = \{drive_B, wait_B\}$ . I.e.,  $Acts = \{drive_A, wait_A, drive_B, wait_B\}$ . To assess

<sup>10</sup> As state above, we assume that the *first* function sorts the action tuples in a deterministic order before returning the first element.

the consequences that include *waiting*, we assume  $B$  is twice as fast as  $A$  (without waiting,  $A$  needs 20 time units to pass the bottleneck while  $B$  needs 10)<sup>11</sup>:

$$a2c_A(acts) := \begin{cases} crash, & \text{if } :acts = (drive_A, drive_B); \\ wait\ 0, & \text{if } :acts = (drive_A, wait_B); \\ wait\ \infty, & \text{if } :acts = (wait_A, wait_B); \\ wait\ 10, & \text{if } :acts = (wait_A, drive_B); \\ null, & \text{otherwise.} \end{cases}$$

$$a2c_B(acts) := \begin{cases} crash, & \text{if } :acts = (drive_A, drive_B); \\ wait\ 20, & \text{if } :acts = (drive_A, wait_B); \\ wait\ \infty, & \text{if } :acts = (wait_A, wait_B); \\ wait\ 0, & \text{if } :acts = (wait_A, drive_B); \\ null, & \text{otherwise.} \end{cases}$$

We construct the following utility quantification functions and subtract an amount proportional to the waiting time from the utility value 1 of *wait 0*:

$$u2c_A(consqs) := \begin{cases} -\infty, & \text{if } :consqs = \{crash\}; \\ 0.9, & \text{if } :consqs = \{wait\ 20\}; \\ 0, & \text{if } :consqs = \{wait\ \infty\}; \\ 1, & \text{if } :consqs = \{wait\ 0\}; \\ null, & \text{otherwise.} \end{cases}$$

$$u2c_B(consqs) := \begin{cases} -\infty, & \text{if } :consqs = \{crash\}; \\ 0.8, & \text{if } :consqs = \{wait\ 20\}; \\ 0, & \text{if } :consqs = \{wait\ \infty\}; \\ 1, & \text{if } :consqs = \{wait\ 0\}; \\ null, & \text{otherwise.} \end{cases}$$

Actions-to-consequences mappings and utility quantification functions can then be combined to utility functions:

$$u_A(acts) = \begin{cases} 1, & \text{if } :acts = (drive_A, wait_B); \\ 0.9, & \text{if } :acts = (wait_A, drive_B); \\ 0, & \text{if } :acts = (wait_A, wait_B); \\ -\infty, & \text{if } :acts = (drive_A, drive_B); \\ null, & \text{otherwise.} \end{cases}$$

$$u_B(acts) = \begin{cases} 0.8, & \text{if } :acts = (drive_A, wait_B); \\ 1, & \text{if } :acts = (wait_A, drive_B); \\ 0, & \text{if } :acts = (wait_A, wait_B); \\ -\infty, & \text{if } :acts = (drive_A, drive_B); \\ null, & \text{otherwise.} \end{cases}$$

<sup>11</sup>  $drive_A \wedge wait_A$  and  $drive_B \wedge wait_B$ , respectively, are mutually exclusive ( $\{drive_A \oplus wait_A, drive_B \oplus wait_B\}$ , with  $A \oplus B := (A \vee B) \wedge \neg(A \wedge B)$ ). I.e., the functions return *null* if  $drive_A \wedge wait_A \vee drive_B \wedge wait_B$ .

We assume scenarios where both agents are *driving* or both agents are *waiting* are not acceptable by either agents and introduce the corresponding acceptability rules:

$$acc_{A,B}(acts) := \begin{cases} false, & \text{if } : acts = (drive_A, drive_B) \vee (wait_A \wedge wait_B); \\ null, & \text{if } : (drive_A \in acts \wedge wait_A \in acts) \vee (drive_B \in acts \wedge wait_B \in acts); \\ true, & \text{otherwise.} \end{cases}$$

Based on the utility functions  $(u_A, u_B)$ , we create new utility functions  $(u'_A, u'_B)$  that consider the acceptability rules:

$$u'_A(acts) := \begin{cases} \frac{1}{2}, & \text{if } : acts = (drive_A, wait_B); \\ \frac{1}{3}, & \text{if } : acts = (wait_A, drive_B); \\ null, & \text{otherwise.} \end{cases}$$

$$u'_B(acts) := \begin{cases} 1, & \text{if } : acts = (wait_A, drive_B); \\ \frac{1}{3}, & \text{if } : acts = (drive_A, wait_B); \\ null, & \text{otherwise.} \end{cases}$$

Finally, we apply the empathic agent algorithms to our scenario. Using the *naive* algorithm, the agents apply the acceptability rules, but do not consider the other agent's strategy. Hence, both agents decide to *drive*, (and consequently *crash*).

$$D\_A\_N_A(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_A) = drive_A$$

$$D\_A\_N_B(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_B) = drive_B$$

The resulting utility is  $-\infty$  for both agents. None of the two other algorithms (*lazy*, *full*) allows any agent to decide to execute an action tuple that does not optimize shared utility. I.e., both algorithms yield the same result:

$$D\_A\_L_A(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_A) = wait_A$$

$$D\_A\_L_B(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_B) = drive_B$$

$$D\_A\_F_A(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_A) = wait_A$$

$$D\_A\_F_B(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_B) = drive_B$$

The resulting utility is 0.9 for agent A and 1 for agent B. As can be seen, the difference between agent types is not always relevant. The following scenario will provide a distinctive outcome for all three agent variants.

### 3.2 Example 2: Concert

As a second example, we introduce the following scenario<sup>12</sup>. Two empathic agents  $\{A, B\}$  plan to attend a concert of music by either *Bach*, *Stravinsky*, or *Mozart* ( $Acts :=$

<sup>12</sup> The scenario is an adjusted and extended version of the ‘‘Bach or Stravinsky? (BoS)’’ example presented by Osborne and Rubinstein [19, p. 15–16]

$\{Bach_A, Stravinsky_A, Mozart_A, Bach_B, Stravinsky_B, Mozart_B\}$ ). A considers the Bach and Mozart concerts of much greater pleasure when attended in company of B (utility of 6, respectively 3) and not alone (either concert: 1). In contrast, the Stravinsky concert yields good utility, even if A attends it alone (4). Attending it in company of B merely gives a utility bonus of 1 (total: 5). B prefers concerts in company of A as well (2 for Stravinsky and 4 for Mozart), but gains little additional utility from attending a Bach concert with A (1.1 with A versus 1 alone) because they dislike listening to A's Bach appraisals. Attending any concert alone yields a utility of 1 for B. As the utility is in this scenario largely derived from the subjective musical taste and social preferences of the agents and to keep the example concise, we skip the actions-to-consequences mapping and construct the utility functions right away<sup>13</sup>:

$$u_A(acts) = \begin{cases} null, & \text{if : } length(acts) \neq 2 \vee \\ & length(set(Bach_A, Stravinsky_A, Mozart_A) \cap set(acts)) \neq 1 \vee \\ & length(set(Bach_B, Stravinsky_B, Mozart_B) \cap set(acts)) \neq 1; \\ 6, & \text{else if : } acts = (Bach_A, Bach_B); \\ 5, & \text{else if : } acts = (Stravinsky_A, Stravinsky_B); \\ 4, & \text{else if : } Stravinsky_A \in acts \wedge Stravinsky_B \notin acts; \\ 3, & \text{else if : } acts = (Mozart_A, Mozart_B); \\ 1, & \text{otherwise.} \end{cases}$$

$$u_B(acts) = \begin{cases} null, & \text{if : } length(acts) \neq 2 \vee \\ & length(set(Bach_A, Stravinsky_A, Mozart_A) \cap set(acts)) \neq 1 \vee \\ & length(set(Bach_B, Stravinsky_B, Mozart_B) \cap set(acts)) \neq 1; \\ 1.1, & \text{else if : } acts = (Bach_A, Bach_B); \\ 2, & \text{else if : } acts = (Stravinsky_A, Stravinsky_B); \\ 4, & \text{else if : } acts = (Mozart_A, Mozart_B); \\ 1, & \text{otherwise.} \end{cases}$$

We introduce the following acceptability function that applies to both agents (although it is of primary importance for agent A). As agent A is banned from the venue that hosts the Stravinsky concert, the action  $Stravinsky_A$  is not acceptable:

$$acc_{A,B}(acts) := \begin{cases} false, & \text{if : } acts = Stravinsky_A \in acts; \\ true, & \text{otherwise.} \end{cases}$$

<sup>13</sup> Note that the if-condition that triggers the return of a *null* value simply defines that  $Bach_A$ ,  $Stravinsky_A$ , and  $Mozart_A$  are mutually exclusive, as are  $Bach_B$ ,  $Stravinsky_B$ , and  $Mozart_B$ .

Considering the acceptability function, we create the following updated utility functions:

$$u'_A(acts) = \begin{cases} \text{null,} & \text{if : } length(acts) \neq 2 \vee \\ & length(set(Bach_A, Stravinsky_A, Mozart_A) \cap set(acts)) \neq 1 \vee \\ & length(set(Bach_B, Stravinsky_B, Mozart_B) \cap set(acts)) \neq 1 \vee \\ & Stravinsky_B \in acts; \\ 6, & \text{else if : } acts = (Bach_A, Bach_B); \\ 4, & \text{else if : } Stravinsky_A \in acts; \\ 3, & \text{else if : } acts = (Mozart_A, Mozart_B); \\ 1, & \text{otherwise.} \end{cases}$$

$$u'_B(acts) = \begin{cases} \text{null,} & \text{if : } length(acts) \neq 2 \vee \\ & length(set(Bach_A, Stravinsky_A, Mozart_A) \cap set(acts)) \neq 1 \vee \\ & length(set(Bach_B, Stravinsky_B, Mozart_B) \cap set(acts)) \neq 1 \vee \\ & Stravinsky_B \in acts; \\ 1.1, & \text{else if : } acts = (Bach_A, Bach_B); \\ 4, & \text{else if : } acts = (Mozart_A, Mozart_B); \\ 1, & \text{otherwise.} \end{cases}$$

Now, we can run the empathic agent algorithms. The *naive* algorithm returns *Bach* for agent *A* and *Mozart* for agent *B*:

$$D\_A\_N_A(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_A) = Bach_A$$

$$D\_A\_N_B(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_B) = Mozart_B$$

The resulting utility is 1 for both agents. The *lazy* algorithm returns *Mozart* for both agents:

$$D\_A\_L_A(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_A) = Mozart_A$$

$$D\_A\_L_B(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_B) = Mozart_B$$

The resulting utility is 3 for agent *A* and 4 for agent *B*. The *full* algorithm returns *Bach* for both agents:

$$D\_A\_F_A(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_A) = Bach_A$$

$$D\_A\_F_B(\{u'_A, u'_B, \}, \{acc_A, acc_B\}, Acts_B) = Bach_B$$

The resulting utility is 6 for agent *A* and 1.1 for agent *B*.

### 3.3 JavaScript Implementation

We implemented the running examples in JavaScript<sup>14</sup>. As a basis for the implementation, we created a simple framework that consists of the following components:

- **Web socket server: environment and communications manager.** The environment and communications interface is implemented by a web socket server that consists of the following components:

<sup>14</sup> The code, as well as documentation and tests, are available at <http://s.cs.umu.se/qxgbfi>.

- **Environment and communications manager.** The web server provides a generic environment and communications manager that relays messages between agents and provides the shared value system of acceptability rules.
  - **Environment specification.** The environment specification contains scenario-specific information and enables the server to determine and propagate the utility rewards to the agents.
- **Web socket clients: empathic agents.** The empathic agents are implemented as web socket clients that interact via the server described above. Each agent consists of the following two components:
- **Generic empathic agent library.** The generic empathic agent library provides a function to create an empathic agent object with the properties *ID*, *utilityMappings*, *acceptabilityRules*, and *type* (*naive*, *lazy*, or *full*). The empathic agent object is then equipped with an action determination function that implements the empathic agent algorithm as described above.
  - **Agent specifications.** The agent specification consists of the scenario-specific information of all agents in the environment, as well as of the current agents' identifier and type (*naive*, *lazy*, or *full*) and is used to instantiate a specific empathic agent. Note that in the implementation, we construct the utility functions right away and do not use actions-to-consequences mappings.

The implementation assumes that the specifications provided to both agents and to the server is consistent. Fig. 1 depicts the architecture of the empathic agent JavaScript implementation for the *vehicle* scenario. We chose JavaScript as the lan-

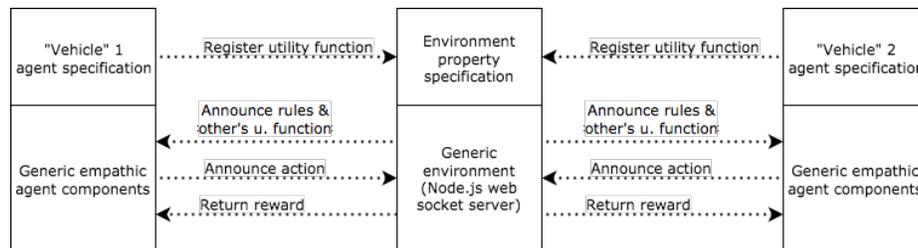


Fig. 1: Empathic intelligent: architecture

guage for implementing the scenario to show how to implement basic empathic agents using a popular general-purpose programming language, but concede that a more powerful implementation in the context of MAS frameworks like Jason is of value.

## 4 Reasoning-loop Architecture

We create a reasoning-loop architecture for the empathic agent and again assume a two-agent scenario to simplify the description. The architecture consists of the following components:

- **Empathic agent (EA).** The empathic agent is the system’s top-level component. It has three generic components (*observer*, *negotiator*, and *interactor*) and five dynamically generated functions/objects (*utility function* and *acceptability function* of both agents, as well as a formalized model of the *shared system of values*).
- **Target agent (TA).** In the simplest scenario, the empathic agent interacts with exactly one other agent (the *target agent*), which is modeled as a black box. Pre-existing knowledge about the target agent can be part of the models the empathic agent has of the target agent’s utility and acceptability functions.
- **Shared system of values.** The shared system of values allows comparing the utility functions of the agents and creating their acceptability functions, as well as their actions-to-consequences mappings and utility quantification functions, from which the utility functions are derived.
- **Utility function.** Based on the actions-to-consequences mappings and utility quantification functions, each empathic agent maintains its own utility function, as well as models of the utility function of the agent it is interacting with.
- **Acceptability function.** Based on the shared system of values, the agent derives the acceptability functions (as described above) to then incorporate them into updated utility functions, which it feeds into the empathic agent algorithm to determine the best possible tuple of actions.
- **Observer.** The observer component scans the environment, registers other agents, *receives* their utility functions, and also keeps the agent’s own functions updated. To construct and update the utility and acceptability functions without explicitly receiving them, the observer could make use of inverse reinforcement learning methods, as for example described by [10].
- **Negotiator.** The negotiator identifies and resolves conflicts of interests using the *acceptability function* models and instructs the interactor to engage with other agents if necessary, in particular, to propose a solution for a conflict of interest, or to resolve the conflict immediately (depending on the level of confidence that the solution is indeed acceptable). The negotiator could make use of argument-based negotiation (see e.g.: [3]).
- **Interactor.** The interactor component interacts with the agent’s environment and in particular with the target agent to work towards the conflict resolution. The means of communication is domain-specific and not covered by the generic architecture.

Fig. 2 presents a simple graphical model of the empathic agent’s reasoning loop architecture.

## 5 Alignment with BDI Architecture and Possible Implementation with Jason

Our architecture reflects the common belief-desire-intention (BDI) model as based on [7] to some extent:

- If a priori available to both agents in the forms of rules or norms, *beliefs*, and *belief sets* are part of the shared value system. Otherwise, they qualify the agents’ utility

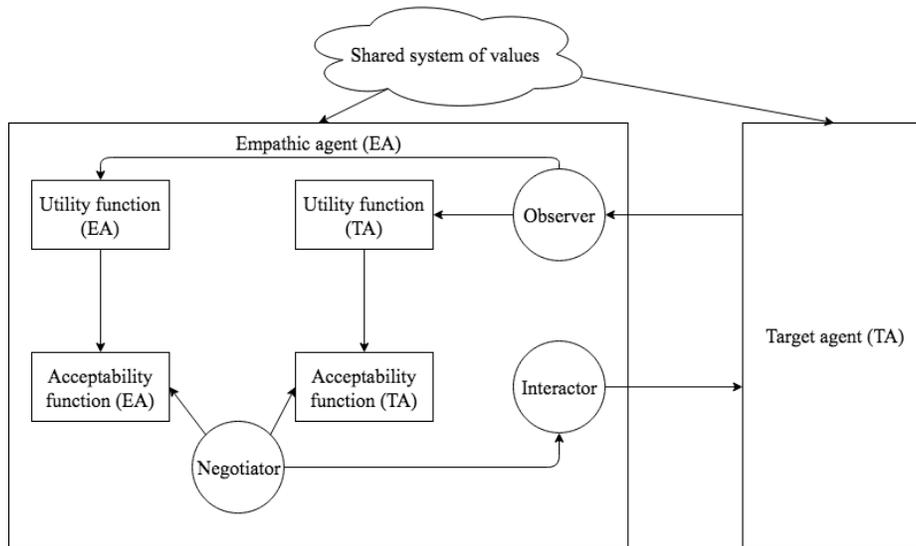


Fig. 2: Empathic intelligent: architecture

and acceptability functions directly. In contrast, *desires* define the objective(s) towards which an agent's utility function is optimized and are—while depending on beliefs—not directly mutable through persuasive argumentation between the agents.

- Intentions are the tuples of actions the agents choose to execute.
- As it strives for simplicity, our architecture does for now not distinguish between desires and goals, and intentions and plans, respectively.

We expect to improve the alignment of our framework with the BDI architecture to facilitate the integration with existing BDI-based theories and implementation using BDI frameworks. The Jason platform for multi-agent system development [6] can serve as the basis for implementing the empathic agent. While simplified running examples of our architecture can be implemented with Jason, extending the platform to provide an empathic agent-specific abstraction layer would better support complex scenarios.

## 6 Discussion

In this section, we place our empathic agent concepts into the context of existing work, highlight potential applications, analyze limitations, and outline future work.

### 6.1 Similar Conflict Resolution Approaches

Our empathic agent can be considered a generic and basic agent model that can draw upon a large body of existing research on multi-agent learning and negotiation techniques for possible extensions. A survey of research on agents that model other agents

is provided by Albrecht and Stone [1]. The idea of combining a utility-based approach with acceptability rules to emulate *empathic* behavior is to our knowledge novel. However, a somewhat similar concept is presented by Black and Atkinson, who propose an argumentation-based approach for an agent that can find agreement with one other agent on acceptable actions and can develop a model of the other agent’s preferences over time [5]. While Black’s and Atkinson’s approach is similar in that it reflects Coplan’s definition of empathy (it maintains “a process through which [it] simulates another’s situated psychological states, while maintaining clear self–other differentiation” [12]) to some extent we identify the following key differences:

- The approach is limited to a two-agent scenario.
- The agent model is preference-based and not utility-based. While this has the advantage that it does not require reducing complex preferences to a simple numeric value, it makes it harder to combine with existing learning concepts (see below).
- The agent has the ability to learn another agent’s preferences over time. However, the learning concept is—according to Black and Atkinson—“not intended to be complete” [5]. We suggest that while our empathic agent does not provide learning capabilities by default, it has the advantage that its utility-based concept allows for integration with established inverse reinforcement learning algorithms (see: Subsection 6.4).
- The agent Black and Atkinson introduce is not *empathic* in that it tries to compromise with the other agent, but rather uses its ability to model the agent’s preferences to improve its persuasive capabilities by tailoring the arguments it provides to this agent.

## 6.2 Potential Real-World Use Cases

In this chapter, we exemplified the empathic agent with two simple scenarios, with the primary purpose of better explaining our agent’s core concepts. These scenarios do not fully reflect real-world use cases. However, the core concepts of the agent can form the basis of solutions for real-world applications. Below, we provide a non-exhaustive list of use case types empathic agents could potentially address:

- **Handling aspects of traffic navigation scenarios that cannot be covered by static rules.** Besides adjusting the assertiveness levels to the preferences of their drivers, as suggested by Sikkenk and Terken [23], and Yusof et al. [26], autonomous vehicles could consider the driving style of other human- or agent-controlled vehicles to improve traffic flow, for example by adjusting speed or lane-changing behavior according to the (perceived) utility functions of all traffic participants or to resolve unexpected incidents (in particular emergencies).
- **Mitigating negative effects of large-scale web applications on their users.** Evidence exists that suggests the well-being of *passive* (mainly content-consuming) users of social media is frequently negatively impacted by technology, while the well-being of at least some users, who actively engage with others through the technology, improves [20]. To facilitate social media use that is positive for the users’ well-being, an empathic agent could serve as a mediator between user needs (social

inclusion) and the business goals of the technology provider (often: maximization of advertisement revenue).

- **Decreasing the negotiation overhead for agent-based manufacturing systems.** Autonomous agent-based manufacturing systems are an emerging alternative to traditional, hierarchically managed control architectures [16]. While agent-based systems are considered to increase the agility of manufacturing processes, one disadvantage of agent-based manufacturing systems is the need for negotiation between agents and the resulting overhead (see for example: Brucocoleri et al. [8]). Employing empathic agents in agent-based manufacturing scenarios can possibly help solve conflicts of interests efficiently.
- **Improving persuasive healthcare technology.** Persuasive technology—“computerized software or information system designed to reinforce, change or shape attitudes or behaviours or both without using coercion or deception” [18]—is frequently applied in healthcare scenarios [11], in particular, to facilitate behavior change. Persuasive functionality is typically implemented using recommender systems [14], which in general struggle to compromise between system provider and end-user needs [21]. This can be considered as a severe limitation in healthcare scenarios, where trade-offs between serving public health needs (optimizing for a low burden on the healthcare system) and empowering patients (allowing for a subjective assessment of health impact, as well as for unhealthy choices to support individual freedom) need to be made. Hence, employing the empathic agent concepts in this context can be considered a promising endeavor.

### 6.3 Limitations

The purpose of this chapter is to introduce *empathic agents* as a general concept. When working towards a practically applicable empathic agent, the following limitations of our work need to be taken into account:

- The agent is designed to act in a fully observable world, which is an unrealistic assumption for real-world use cases. For better applicability, the agent needs to support probabilistic models of the environment, the other agents, and the shared value system.
- Our formal empathic agent description is logic-based. Integrating it with Markov decision process-based inverse reinforcement learning approaches is a non-trivial endeavor, although certainly possible.
- In the example scenarios we provided, all agents are identically implemented empathic agents. An empathic agent that interacts with non-empathic agents will need to take into account further game-theoretic considerations and to have negotiation capabilities.
- The presented empathic agent concepts use a simple numeric value to represent the utility an agent receives as a consequence of the execution of an action tuple. While this approach is commonly employed when designing utility-based autonomous agents, it is an oversimplification that can potentially limit the applicability of the agent.

- Software engineering and technological aspects of empathic agents need to be further investigated. In particular, the implementation of an empathic agent library using a higher-level framework for multi-agent system development, as we discuss in Section 5 could provide a more powerful engineering framework for empathic agents.

## 6.4 Future Work

We suggest the following research to address the limitations presented in Subsection 6.3:

- So far, we have chosen a logic-based approach to the problem in focus to allow for a minimalistic problem description with low complexity. Alternatively, the problem could be approached from a *reinforcement learning* perspective (see for an overview of multi-agent reinforcement learning: [9]). Using (partially observable) Markov decision processes, one can introduce a well-established temporal and probabilistic perspective<sup>15</sup>. A key capability our empathic agent needs to have is the ability to learn the utility function of other agents. A comprehensive body of research on enabling this ability by applying *inverse reinforcement learning* exists (for example: [10] and [17]). Hence, creating a Markovian perspective on the empathic agent to enable the application of reinforcement learning methods for the observational learning of the utility functions of other agents can be considered relevant future work.
- To better assess the applicability of the empathic agent algorithms, it is important to analyze its computational complexity in general, as well as to evaluate it in the context of specific use cases that might allow for performance-improving adjustments.
- To enable empathic agents to reach consensus in case of inconsistent beliefs argumentation-based negotiation approaches can be applied that consider uncertainty and subjectivity (e.g. [15]) for creating solvers for finding compromises between utility/acceptability functions. Similar approaches can be used to enhance utility quantification capabilities by considering preferences and probabilistic beliefs.
- The design intention of the architectural framework we present in Section 4 is to form a high-level abstraction of an empathic agent that is to some extent agnostic of the concepts the different components implement. We are confident that the framework can be applied in combination with existing technologies to create a real-world applicable empathic agent framework, at least for use cases that allow making some assumptions regarding the interaction context and protocol.
- The ultimate goal of this research is to apply the concept in a real-world scenario and evaluate to what extent the application of empathic agents provides practically relevant benefits.

---

<sup>15</sup> However, the same can be achieved with temporal and probabilistic logic.

## 7 Conclusion

In this chapter, we introduced the concept of an *empathic agent* that proactively identifies potential conflicts of interests in interactions with other agents and uses a mixed utility-based/rule-based approach to find a mutually acceptable solution. The theoretical framework can serve as a general purpose model, from which advanced implementations can be derived to develop socially intelligent systems that consider other agents' (and ultimately humans') welfare when interacting with their environment. The example implementation, the reasoning-loop architecture we introduced for our empathic agent, and the discussion of how the agent can be implemented with a belief-desire-intention approach provide first insights into how a more generally capable empathic agent can be constructed. As the most important future research steps to advance the empathic agent, we regard the conceptualization and implementation of an empathic agent with learning capabilities, as well as the development of a first simple empathic agent that solves a particular real-world problem.

**Acknowledgements** We thank the anonymous reviewers for their constructive critical feedback. This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

## References

1. Albrecht, S.V., Stone, P.: Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence* **258**, 66–95 (May 2018)
2. Alshabi, W., Ramaswamy, S., Itmi, M., Abdulrab, H.: Coordination, cooperation and conflict resolution in multi-agent systems. In: Sobh, T. (ed.) *Innovations and Advanced Techniques in Computer and Information Sciences and Engineering*. pp. 495–500. Springer Netherlands, Dordrecht (2007)
3. Amgoud, L., Dimopoulos, Y., Moraitis, P.: A unified and general framework for argumentation-based negotiation. In: *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*. pp. 158:1–158:8. AAMAS '07, ACM, New York, NY, USA (2007)
4. Berinsky, A.J.: Rumors and health care reform: experiments in political misinformation. *British Journal of Political Science* **47**(2), 241–262 (2017)
5. Black, E., Atkinson, K.: Choosing persuasive arguments for action. In: *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*. pp. 905–912. International Foundation for Autonomous Agents and Multiagent Systems (2011)
6. Bordini, R.H., Hübner, J.F.: BDI agent programming in AgentSpeak using Jason. In: *International Workshop on Computational Logic in Multi-Agent Systems*. pp. 143–164. Springer (2005)
7. Bratman, M.: *Intention, Plans, and Practical Reason*. Center for the Study of Language and Information (1987)
8. Bruccoleri, M., Nigro, G.L., Perrone, G., Renna, P., Diega, S.N.L.: Production planning in reconfigurable enterprises and reconfigurable production systems. *CIRP Annals* **54**(1), 433–436 (2005)

9. Busoniu, L., Babuska, R., De Schutter, B.: A comprehensive survey of multiagent reinforcement learning. *IEEE Trans. Systems, Man, and Cybernetics, Part C* **38**(2), 156–172 (2008)
10. Chajewska, U., Koller, D., Ormoneit, D.: Learning an agent’s utility function by observing behavior. In: *ICML*. pp. 35–42 (2001)
11. Conroy, D.E., Yang, C.H., Maher, J.P.: Behavior change techniques in top-ranked mobile apps for physical activity. *American journal of preventive medicine* **46**(6), 649–652 (2014)
12. Coplan, A.: Will the real empathy please stand up? a case for a narrow conceptualization. *The Southern Journal of Philosophy* **49**(s1), 40–65 (2011)
13. Dautenhahn, K.: The art of designing socially intelligent agents: Science, fiction, and the human in the loop. *Applied artificial intelligence* **12**(7-8), 573–617 (1998)
14. Hors-Fraile, S., Rivera-Romero, O., Schneider, F., Fernandez-Luque, L., Luna-Perejon, F., Civit-Balcells, A., de Vries, H.: Analyzing recommender systems for health promotion using a multidisciplinary taxonomy: A scoping review. *International journal of medical informatics* **114**, 143–155 (2018)
15. Marey, O., Bentahar, J., Khosrowshahi-Asl, E., Sultan, K., Dssouli, R.: Decision making under subjective uncertainty in argumentation-based agent negotiation. *Journal of Ambient Intelligence and Humanized Computing* **6**(3), 307–323 (Jun 2015)
16. Monostori, L., Váncza, J., Kumara, S.: Agent-based systems for manufacturing. *CIRP Annals* **55**(2), 697 – 720 (2006)
17. Ng, A.Y., Russell, S.J., et al.: Algorithms for inverse reinforcement learning. In: *Icml*. pp. 663–670 (2000)
18. Oinas-Kukkonen, H., Harjumaa, M.: Towards deeper understanding of persuasion in software and information systems. In: *Advances in computer-human interaction, 2008 first international conference on*. pp. 200–205. *Ieee* (2008)
19. Osborne, M.J., Rubinstein, A.: *A course in game theory*. MIT press (1994)
20. Philippe, V., Oscar, Y., Maxime, R., John, J., Ethan, K.: Do social network sites enhance or undermine subjective well-being? a critical review. *Social Issues and Policy Review* **11**(1), 274–302
21. Ricci, F., Rokach, L., Shapira, B.: Recommender systems: introduction and challenges. In: *Recommender systems handbook*, pp. 1–34. Springer (2015)
22. Russell, S.J., Norvig, P.: *Artificial intelligence: a modern approach*. Malaysia; Pearson Education Limited, (2016)
23. Sikkenk, M., Terken, J.: Rules of conduct for autonomous vehicles. In: *Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. pp. 19–22. *AutomotiveUI ’15*, ACM, New York, NY, USA (2015)
24. Von Neumann, J., Morgenstern, O.: *Theory of games and economic behavior*. *Bull. Amer. Math. Soc* **51**(7), 498–504 (1945)
25. Wojdyski, B.W., Bang, H.: Distraction effects of contextual advertising on online news processing: an eye-tracking study. *Behaviour & Information Technology* **35**(8), 654–664 (2016)
26. Yusof, N.M., Karjanto, J., Terken, J., Delbressine, F., Hassan, M.Z., Rauterberg, M.: The exploration of autonomous vehicle driving styles: Preferred longitudinal, lateral, and vertical accelerations. In: *Proceedings of the 8th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. pp. 245–252. *Automotive’UI 16*, ACM, New York, NY, USA (2016)