

la VIDA: Towards a Motivated Goal Reasoning Agent

Ursula Addison

uaddison@gradcenter.cuny.edu

Computer Science, CUNY Graduate Center, 365 5th Ave, New York, 10016, NY, USA

Abstract—An autonomous agent deployed to operate over extended horizons in uncertain environments will encounter situations for which it was not designed. A class of these situations involves an invalidation of agent goals and limited guidance in establishing a new set of goals to pursue. An agent will benefit from some mechanism that will allow it to pursue new goals under these circumstances such that the goals are broadly useful in its environment and take advantage of its existing skills while aligning with societal norms. We propose augmenting a goal reasoning agent, i.e., an agent that can deliberate on and self-select its goals, with a motivation system that can be used to both constrain and motivate agent behavior. A human-like motivation system coupled with a goal-self concordant selection technique allows the approach to be framed as an optimization problem in which the agent selects goals that have high utility while simultaneously in harmony with its motivations. Over the agent’s operational lifespan its motivation system adjusts incrementally to more closely reflect the reality of its goal reasoning and goal pursuit experiences. Experiments performed with an ablation testing technique comparing the average utility of goals achieved in the presence and absence of a motivation system suggest that the motivated version of the system leads to pursuing more useful goals than the baseline.

I. INTRODUCTION

An autonomous agent which operates over long-term horizons will encounter situations for which it was not specifically designed; uncertainty and surprises are unavoidable features of complex environments. The impact of this uncertainty can vary widely depending on the specifics of the situation. It’s easy to imagine that the impact could be extreme and will need to be mitigated if autonomous agents are to operate successfully in less controlled settings. We consider the specific scenario where uncertainty arises from a lack of external input or guidance from a controlling entity. If the agent is receiving limited external direction, but is otherwise fully functional, it may adopt one or more policies to determine its next steps. The agent may stop operating or continue taking actions, but in the latter case taking action without a mode designed for increased self-direction may result in various issues. In the most positive and benign outcomes the result could be very inconvenient, prohibitively expensive, and impractical. The space of potential problems is massive ranging from innocuous misunderstandings, to inappropriate or low utility actions, and so on.

We believe this is a worthwhile problem to explore and that it will be impossible to have truly autonomous agents if they are unable to appropriately prepare for and adapt in such situations. Goal reasoning, the ability to self-select and identify

goals, is one solution to this problem. However, goal reasoning implementations that do not allow for behavior that aligns with agent internal motivations, values and experiences may not be sufficiently flexible for operation over indefinite periods in the proverbial “jungle” that characterizes real-world environments [Addison, 2023]. We propose augmenting the typical externally motivated goal reasoning agent with a motivation system that is continuously refined by the agent’s experiences. In this way, the agent will have a value and identity system that allows it to interact with its environment in a fitting and responsible manner, while the system is incrementally adjusted over time to more closely reflect the realities of the agent’s experiences.

We pose this problem as the two research questions:

- 1) How can an artificial agent utilize its identity and value system to self-select, manage, and realize its objectives?
- 2) How can the agent’s experiences of self-selecting, managing, and realizing its objectives be incorporated into its identity and value system to influence future behavior?

From these research questions we can extract several sub-problems, which we attempt to address throughout the article while comparing and contrasting our solution to other solutions with the same aims. We must first determine a model for the agent’s identity, value system, and internal state that supports their representation and provides a set of processes that manage these elements and their impact on self-selecting, managing, and realizing objectives. There additionally needs to be some formalism or technique that can explicitly represent and manage these objectives, i.e., goals. A technique to select and achieve goals on its own is not sufficient, there needs to be some criteria or mechanism that assures that the selected objectives are in alignment with the agent’s identity and value system; there currently exists many systems that address some of our stated sub-problems [Gajderowicz et al., 2018] [Coman and Muñoz-Avila, 2014] [Swoboda et al., 2022] [Samsonovich, 2013] [Yu et al., 2021] [Sun, 2009], however, to our knowledge the use of the empirically based goal-self concordance theory to achieve agent-goal alignment [Sheldon and Elliot, 1999] [Sheldon, 2014] [Milyavskaya et al., 2014] is novel. And lastly, we will also need to establish how the results of managing and achieving goals impacts the identity and value system. These sub-problems do not fully resolve our research questions and the related issues, but they are the core problems we will be attempting to address with this paper.

This article is organized as follows: Section 2 offers a

brief coverage of important background information necessary to appreciate the problem and suggested solution. Section 3 acknowledges related work and its connections to the presented work. In Section 4 we introduce our proposed method and some implementation details. In Section 5 we outline our experimental approach and results and the article closes in Section 6 with some conclusions and future research directions.

A. A Motivating Example

A life-long, autonomous agent outfitted with a motive management framework, built and deployed as part of the Alfred, *Home Help Bot Series* [Hawes, 2011] is separated from its client family during a voyage by sea. After the separation, the agent finds itself on a completely unfamiliar landmass utterly at a loss for what goal it should pursue. There is no garbage to be collected, no clothing to wash and store, or no floor to vacuum. However, aside from the seemingly monumental issues of a completely invalidated set of goals and loss of contact with its client family, the agent detects no other malfunction. With a fully charged battery and all its functionality intact what should it do? Should it deactivate itself? Or is there a means for which it can continue having a purpose from which it can generate intelligent, useful, and appropriate goal directed behavior?

The agent is a home butler in a future time when the necessary technology is sufficiently advanced and trustworthy to be deployed in a typical family home. As such, one can imagine that the agent is highly skilled and possesses some abilities near or beyond the human level. It seems feasible that the agent would be able to perform tasks outside of the domain of robotic house servant. Beyond its physical capabilities, the agent’s motivation management framework allows it to manage its goals, communicate effectively with its owners, anticipate their needs, express emotions and so on. In its role as home help bot, the agent is typically allowed to independently manage an assigned set of goals and determine how and when they are achieved. Personality wise, Alfred is a helpful, friendly, and meticulous home help bot that can blend into the background and complete its tasks. We envision this domain which we call the *Cast Away*, the problem of needing new goals, and solution as taking place in a virtual environment. Throughout this article we will refer to Alfred whenever an illustrative example is needed, while following Alfred’s trajectory from a stable, established family home environment to a desolate and unexplored island.

II. BACKGROUND

We have opted for an adaptive agent procedure emphasizing human psychology in lieu of an autonomic engineering heavy approach. From this perspective, our problem and solution are interdisciplinary and incorporate ideas from a variety of fields. For each of the sub-problems identified in the previous section, we model a human trait or use an existing artificial intelligence (AI) technique. To address the sub-problem to represent and manage the agent’s identity, value system, and internal state we

use ideas from Personality Systems Interaction (PSI) Theory [Kuhl and Baumann, 2021] to create a motivation process that manipulates the agent’s self-system, i.e., identity, value system, and internal state which are modeled on the Five Factor Model (FFM) [Howard and Howard, 1995] and Schwartz’s 10 value theory [Schwartz, 2012]. The Personality System Interaction Theory (PSI) developed by Kuhl [Kuhl, 2000] is a motivational model that employs two major systems, the self and volitional facilitation systems; these systems operate on the agent’s self-system to result in a motivation process. The self facilitation system is made of two subsystems, the *object recognition system* (ORS) and the *extension memory* (EM). The volitional facilitation system is formed by two subsystems, the low-level *intuitive behavior control* (IBC) system and the high-level *intention memory* (IM). While we do not directly include these subsystems in la VIDA, each one corresponds to a combination of motivation constructs and processes.

Howard et. al and Schwartz have created two theories to describe personality and a value system respectively with a minimal set of factors. The Five-Factor Model (FFM) is a personality paradigm that uses five dimensions to describe a wide variety of personality configurations [Howard and Howard, 1995]. The CANOE factors are: Conscientiousness, Agreeableness, Neuroticism, Openness, and Extraversion. The *Schwartz theory of basic values* [Schwartz, 2012] is a set of 10 values thought to be present among humans irrespective of the time and place. Each value expresses priorities and motivations and have positive and negative correlations with other values. The ten values are: self-direction, stimulation, hedonism, achievement, power, security, conformity, tradition, benevolence, and universalism.

Goal reasoning in artificially intelligent systems and self-regulation in humans are volitional mechanisms that have the end of selecting goals and determining how and when they are pursued [Aha, 2018] [Cox, 2017] [Mann et al., 2013] [Schank and Abelson, 1977] [Sheldon and Elliot, 1999]. Goal pursuit and achievement among humans is described by Self-Regulation Theory [Mann et al., 2013] which maps very well onto the AI technique called goal reasoning (GR) [Aha, 2018][Cox, 2017][Muñoz-Avila, 2018][Klenk, 2010], which is our selected formalism to address the sub-problem to represent and manage agent objectives; GR however, doesn’t incorporate the motivation system effects experienced by humans during goal pursuit activities, so we model a subset of those described by self-regulation. Motivation system effects resulting from goal pursuit are summarized and used to calibrate agent motivation constructs to impact future goal pursuit.

The Self-Concordance Model (SCM) describes what processes are involved in the human selection of self-aligning goals and how it impacts their motivation systems; we include some motivation elements relevant to SCM into the agent’s motivational constructs and manipulate them according SCM processes to solve our final problem of allowing the agent to select self-concordant goals.

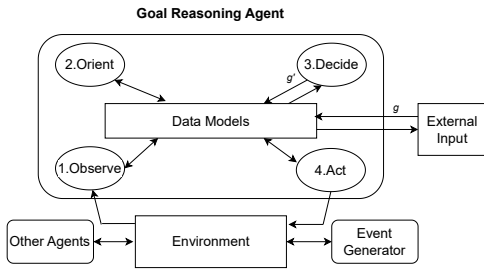


Fig. 1: A general model of a goal reasoning agent, with g' as a formulated goal and g as an assigned goal [Aha, 2018] Boyd et al. [2018].

III. RELATED WORK

This research equips a goal reasoning agent with a simple motivation system, such that its goal reasoning is guided by intrinsic and identified motivations and their related processes with the end of creating a motivated goal reasoning system. An intrinsic motivation results in agent enjoyment and satisfaction while an identified motivation emerges from the values and beliefs of the agent. In this section, we compare and contrast our work with other goal reasoning systems that also model intrinsic or identified factors following a goal reasoning systems survey [Addison, 2023]. BRAMA [Gajderowicz et al., 2017] is a goal reasoning system that models agent motivations according to Maslow’s hierarchy of needs and manages a set of predefined agent-specific goals. Like BRAMA, la VIDA also has a motivation system, but it is based on PSI, FFM, and Schwartz’ 10 values and can reason over arbitrary sets of goals. The motivated rebel agents framework [Coman and Muñoz-Avila, 2014] is based on the goal-driven autonomy (GDA) paradigm where the agent responds to motivation discrepancies by changing its goals. Our work also includes motivations and goal formulation triggers, a key difference being that motivated rebel agents seek self-goal alignment above all else, where as our agent attempts to balance value alignment with goal utility. CLIPS Executive [Swoboda et al., 2022] is a goal reasoning system, that like our system is based on Roberts et al. [Roberts et al., 2016] GR formalism. It also implements GTN planning and the goal reasoning lifecycle as our approach does. A key difference between our work however, is that CLIPS Executive focuses on multi-agent systems of similar, collaborating, agents, where our system only considers a single and highly idiosyncratic agent. eBICA [Samsonovich, 2013] is an architecture created in response to the need for humans and artificial agents to relate on a social and emotional level; it is similar to our work in that eBICA agent behavior is impacted by affect. EDA [Yu et al., 2021] is an architecture that uses well-being as a metric and adapts agent behavior through the reinforcement of successful goal reasoning policies. Our approach also incorporates well-being and allows the agent to adapt, but our well-being metric is a function of goal-self concordance, goal utility, and agent rewards where EDA well-being is an accumulation of positive and negative state changes. CLARION incorporates both

emotions and motivations and is a dual process architecture with implicit and explicit processes [Sun, 2009]. While it extensively implements dual process theory, we only apply dual process concepts in the categorization and accessibility of motivation system components. SOAR Laird [2012] is a general cognitive architecture that attempts to approximate complete rationality using available knowledge while operating. This architecture is equipped with the primitive cognitive building blocks to exhibit much more complicated intelligent behavior, task planning and execution, knowledge representation, and interaction with its environment. ACT-R Anderson [2002] is a cognitive architecture and model of the human mind. Its organization as a collection of modules mimics that of the human brain, with each module producing some aspect of cognition. Sigma Rosenbloom et al. [2017] is a cognitive architecture that combines past work on cognitive architectures with graphical models. Four tenants of its development are: grand unification, generic cognition, functional elegance, and sufficient efficiency.

IV. MATERIALS AND METHODS

In this section we detail the design and implementation of each of the system’s components, organizing our discussion according to the system overview diagram in figure 4.

A. World Knowledge

We currently make the assumption that our agent Alfred will have a high-level of intelligence, capability, and trustworthiness, suitable to be deployed with limited supervision as a home help bot or other role in various settings. This assumption creates a need for the agent to have common sense knowledge about the world as well as what goals can be pursued and how to achieve them. For testing and experimentation purposes we create a toy set of world knowledge that serves as a placeholder for what we envision will be a much larger set in future work. World knowledge about goals, the goal pursuit context, and motivation semantics is stored in the goal library, plan library, and common sense classes.

1) *Goal Library*: The goal library is a class that contains the set of possible goals, and is organized as the core parts: `goal_case_library`, `goal_state_space`, and `goal_library`. The `goal_case_library` is a container that holds goal cases. Each goal class, i.e., a more general goal structure that has not yet been grounded, corresponds to a single goal case which holds all goal class information except its goal state. All state objects, whether the world state or goal states, are a collection of key-value pairs for a set of properties; if the property is *clean* key-value pairs could be (Alfred, true) and (kitchen, false). The numeric ranges of goal case elements were chosen to put boundaries on the calculations performed with them, e.g., total need satisfaction should not exceed 1, thus the four need satisfaction components have an upper bound of 0.25. Following is a list of the data in a goal case with its description; unless a function exists to calculate the element or it is stated otherwise, the data is currently hard coded:

- 1) need satisfaction - partial satisfaction for one or more of the need satisfaction components autonomy, power, affiliation, and achievement $\in [0, 0.25]$
- 2) difficulty - subjective assessment of how easy the goal will be to achieve $\in \{1, \dots, 5\}$
- 3) urgency - reflects time pressure to achieve the goal $\in [0, 5]$, it increases a fraction of its value for each cycle since the goal was last pursued and resets to 0 after being selected
- 4) utility - measure of the goal's positive impact on the agent and environment $\in \{0, \dots, 20\}$
- 5) domain need satisfaction - domains in which need satisfaction is enhanced with their enhancement value $\in [0.0.25]$
- 6) personality - values for each of the FFM 5 factors $\in [0, 1]$
- 7) Schwartz values - entry for 10 universal values $\in [-1, 1]$
- 8) drive - a list of associated drives
- 9) preconditions - a list of state properties that must hold in the environment

The `goal_state_space` holds all goal states that correspond with a goal class and its goal case. The `goal_library` is a container that holds a set of `GoalNodes`, the goal structure that is manipulated by the GR process. Each `GoalNode` is created and loaded with its matching goal state from the `goal_state_space` and goal information from its associated goal case in the `goal_case_library`. The design of our `GoalNode` was developed by Alford et al. [2016] and is detailed in the GR system section. The `filter_goal_library` function uses preconditions and the world state to filter out any possible goals for which the basic preconditions don't hold, e.g., if Alfred wanted to cook and there are no food objects, that goal would be filtered out. The `initialize_goal_library` function brings together all the elements described in this subsection to load the `goal_library` that will later be ingested by the GR process.

2) *Plan Library*: To simplify this iteration of the system, we have opted to use a plan library which contains plan recipes for each of the goal classes in the goal library. Consistent with Alfred's roles, abilities, and environment we implement the plans for possible goals in table III with the data members and member functions in the `PlanLibrary` class. Our `GoalNode` structure has a goal data member that is a goal-task network (GTN), this impacts planning since a `goal_state_space` goal, i.e., root goal, must be decomposed into its subgoals which are then decomposed into a combination of goal methods, task methods, and operators; the GTN plan terminates as a sequence of operators which are then executed. Each member function of the `PlanLibrary`, except for the constructor, is a task method (if appended with `m_`) or operator. Methods select the appropriate operator which can directly modify its argument's state during execution; each of the `PlanLibrary` functions is described in table II.

3) *Common Sense*: This class is intended to hold common sense information that wouldn't easily be stored in the other data structures. Presently, its only role is to map Alfred's roles

to a set of drives. The `role_set_to_drive_set` function compares the roles stored in Alfred's identity profile to a set of drives known to the system, if the conditions hold for a drive, it will be added to Alfred's drive set. Commonsense knowledge is manually encoded in map containers where each element is a key-value pair. For example, Alfred has the role *housekeeper*, so the conditions for the drive *maintain order* will hold and be added to the drive set; drives may be conflicting, we don't limit the combinations to those that are in harmony, when drives conflict the drive with the greater strength will dominate in determining behavior.

B. GR System

A goal reasoning system is one that can reason over and self-select its goals Aha [2018]. Specifically, a GR system must meet a set of minimum requirements: explicit goal representation, formulation of new goals in response to a trigger, and management of formulated goals Vattam et al.. In this section, we outline the GR system component of la VIDA, providing details for each of its subcomponents. Alford et al. [2016] proposed a formalism for hierarchical planning with networks that contain goals, tasks, and operators called goal-task networks (GTN). Roberts et al. [2016] use this goal representation in their formalism and we follow suit. There are some slight differences as our goals are conceptually designed as GTNs but implemented as functions similar to work by Nau et. al with the GTPyhop planner Nau et al. [2021]. A GTN will contain a partially ordered set of goals, subgoals, composite tasks, primitive tasks, and operators; to execute an expansion of a goal node, goals and composite tasks should be decomposed. We preserve the order relation between goals and subgoals using the goal tree structure employed by CLIPS Executive. Having discussed a goal object, we can now describe a goal node object which is the key data structure manipulated by a GR process. Our goal node is implemented following work by Alford et al. [2016] and has data members as follows:

- 1) *g* - a goal object
- 2) *C* - a set of 0 or more goal constraints or preconditions
- 3) *o* - goal life cycle status
- 4) *X* - set of all goal plans
- 5) *x* - selected goal plan
- 6) *q* - set of goal quality metrics, it currently only holds inertia which is incremented by 1 on every status change

1) *GR Process*: The core component of the system is its goal reasoning process. We adapt our GR process from the GR process developed for the ACTORSIM framework Roberts et al. [2016] which progresses goal nodes, i.e., the data structure which contains a goal, through goal memory via the application of refinement strategies. Once formulated, each goal has a status which indicates its place in the goal life cycle and determines what refinement strategies can be applied to it. A GR process has the four phases goal formulation, goal management, planning, and execution; a subset of refinement strategies implicitly form each of the GR phases. Except for the ACTIVATE, MONITOR, RESOLVE-BY, FAIL-TO, and

PROCESS strategies, a refinement strategy applied to a goal node will update its status to the strategy name.

Algorithm 1 Baseline GR Process

Input: Adopted goal set size, goal memory size, n GR cycle count, output file
Output: GR cycle output for m goals
initialize adopted goal set;
for n goal reasoning cycles **do**
 rank goal library;
 update goal memory to FORMULATE, DROP, or re-prioritize goal nodes;
 update adopted goal set to SELECT the top m goals;
 ACTIVATE top goal node;
 EXPAND active goal node;
 COMMIT shortest expansion;
 DISPATCH active goal node;
 if outcome = SUCCESS **then**
 FINISH active goal node;
 else if outcome = FAILURE **then**
 EVALUATE active goal node;
 PROCESS event;
 end if
 update urgency for all goals;
 clean up active goal node;
 result := goal utility, outcome, name, plan, updated world state;
end for
return goal reasoning process output for m goals

FORMULATE, SELECT, and ACTIVATE comprise the goal formulation phase, where Alfred modifies its adopted goal set, i.e., the set of goals marked for goal pursuit. FORMULATE initializes a goal node object and adds it to goal memory. SELECT updates a formulated goal node's status from FORMULATED indicating a goal node simply lives in goal memory to SELECTED, indicating that it can pursued. ACTIVATE points to the highest ranked goal node in goal memory as the current goal to pursue. Goal management phase strategies include: RESOLVE-BY-DEFER, RESOLVE-BY-REFORMULATE, DROP, PROCESS, FAIL-TO, and FINISH strategies. RESOLVE-BY-DEFER and RESOLVE-BY-REFORMULATE will set an EVALUATED goal node to SELECTED and FORMULATED respectively. A goal node which has been processed and no longer has any subgoals is marked FINISHED by FINISH. A goal node may be erased from goal memory by passing it to DROP. Any goal may be regressed to its previous status by invoking FAIL-TO on it. PROCESS may be called on any goal node at anytime and acts as a messenger to bring information to goal nodes about the GR process, e.g., success or failure status.

The strategies EXPAND and COMMIT form the planning phase, while DISPATCH, MONITOR, and EVALUATE form the execution phase. EXPAND decomposes the active goal node into subgoals and plans for each subgoal. This plan is appended to the goal node's X data member. COMMIT sets the goal node's x data member to the first expansion in X , as the plan that will be executed to achieve the goal. DISPATCH then executes each subgoal plan in the selected expansion which decreases Alfred's energy level a small amount. The

goal reasoning process may be invoked by the MONITOR strategy, that catches any anomalies that occur while the expansion is being carried out in the world state; if an anomaly, thrown as a *gr_process_exception* by the system, is caught, EVALUATE is applied to signal that an event has occurred. In an EVALUATED goal node state, our PROCESS strategy which also handles thrown *gr_process_exceptions* is tasked with graceful recovery or termination of the current goal reasoning cycle. Output for a run of GR process will report:

- the **active goal** which corresponds to the goal class name for the active goal node
- the goal pursuit **outcome** as SUCCESS or FAILURE
- a **plan** as a collection of subgoal plans
- the **world state**
- the amount of **utility** Alfred garnered

2) *Formulation Triggers*: Currently goal node formulation triggers are fired when the adopted goal set or goal memory drops beneath the sizes specified as part of the GRProcess object's instantiation. As noted earlier, goal formulation involves the identification and selection of the next best goal to pursue, instead of actually forming new goals, in large part, we believe that the repository of common knowledge would include a great number of possible goals to pursue and could include most of the goals Alfred might like to pursue in variety of environments.

C. GR Agent Motivation System

The motivation system is modeled on PSI theory and configured with FFM and Schwartz theory of basic values elements. Two novel constructs which we have termed an *identity profile* Addison [2022] and *motive network* make up the core explicit representation of the agent's motivation system. An identity profile is the self-system construct used to characterize the motive configuration of an agent, while the motive network encapsulates the agent's goal pursuit experiences and memories. We have chosen to construct our identity profile following work by FFM Howard and Howard [1995] and Schwartz' 10 values because both theories use a small set of factors to express a personality or value system, supporting an expansive range of identities.

We implicitly implement PSI in la VIDA using combinations of constructs and subprocesses. EM is captured in the identity profile and motive network, while ORS calculates the need satisfaction and expectancy-value of the active goal node; expectancy-value is the agent's best guess at the "worth" of a goal, i.e., the product of goal importance with the ratio of volitional competence to goal difficulty. IM is represented by the agent's adopted goal set, conation variables, and the formulation, management, and planning phases. IBC is also represented by the planning phase and additionally the execution phase. System interactions are facilitated by changes in the agent's affect level; the agent can experience negative, negative down-regulated, neutral, positive down-regulated, and positive affect. Different transitions between affect levels are responsible for the emergence and retreat of specific capabilities and subprocess. The full implementation

of this motivation system, specifically affective facilitation, is future work, presently affect is only utilized in the incremental decrease of agent well-being and is otherwise only tracked and modified by the system.

1) *Identity Profile*: The identity, which is an important element of an agent's self-system, is a primary source of motivation Sheldon [2014]. Our identity profile, which models the identity, is composed of an agent's values, roles, beliefs, and personality and is realized by the IdentityProfile class. The role set holds agent roles as natural language, personality holds the FFM factors, and the value system holds the 10 values from *Schwartz theory of basic values*; personality and value system elements may be assigned a numerical value in the intervals [0, 1] and [-1, 1] respectively with lower numbers linked to lower strength of the value and higher numbers to higher strength. Agent beliefs, are states that the agent holds to be true, for example, a home help bot may hold the belief that the house owner is the boss, expressed as the key-value pair (house owner, true) for the *boss* property. The need satisfaction holds the optimal need satisfaction amounts the agent needs to receive for the need satisfaction components autonomy, achievement, power, and affiliation. The need domain holds a list of domains for which the agent receives greater need satisfaction irrespective of the goal class it is pursuing, need domains are represented with natural language. Similarly, the drive set holds a set of the agent's drives in natural language; upon initialization the drive set is empty and is updated by using world knowledge information such as possible drives. Volitional competence is a weight that can be used to add a probabilistic nature to an agent's motivation maintenance competencies, i.e. callable conation variables, that become successful or unsuccessful with some probability. Conation variables are described in an upcoming subsection.

2) *Motive Networks*: A motive network is a structure we use to record qualitative information about the agent's goal pursuit experiences; its inception is inspired by human autobiographical memories Heckhausen and Heckhausen [2008]. Every goal class that has been pursued by Alfred will be represented in the motive network by a motive node, which is updated each time the goal class is pursued. Information tracked by a motive node includes: the goal class name, goal success count, goal failure count, average well-being, average effort, average frustration, the last affective state, the max well-being ever experienced for the goal class and its corresponding world state, and finally the motive disposition which captures the motivation profile of a goal class.

3) *Conation Variables*: Next we detail the design and semantics of a suit of conation variable functions and related utility functions. Previously we discussed the identity and processes involved in its management, some of these processes can be harnessed by personality variables and characteristic behavior which play a key role in setting self-concordant goals. They facilitate communication between the type 2 explicit goal setting system and type 1 implicit memory, affective, conative, and cognitive processes. Signals and feedback from type 1 systems may be experienced as feelings which are

more easily perceived when the agent feels positive affect and embraces introspection, mindfulness, intuition, and creativity. Conation variables are *intuition* (reveals information in the identity profile), *self-relaxation* (manages negative affect), *self-motivation* (manages positive affect), and *self-perception* (access goal-self concordance outputs).

4) *Personality Orientation*: An agent's motivational configuration may skew towards an action or state orientation; agents that are action oriented emphasize taking action and agents that are state oriented emphasize state information and past occurrences. In la VIDA agent personality orientation impacts other motivation system elements, the stronger orientation, i.e., state or action is first determined and there after that orientation will influence a set of motivation components. An action oriented agent's extraversion factor is a function of the action orientation value and as a result, it indirectly acts as a weight on the competency to self-motivate, i.e., the conation variable *self-motivation*. An action oriented agent will also have a tendency to hyper-focus on positive outcomes resulting from self-regulation. In addition the agent is impacted more strongly by negative affect than an state oriented agent, thus the negative affect threshold to activate the IM, that is the formulation and planning phases, is lowered; conversely the positive affect threshold required to activate IBC, that is the execution phase, is increased. A state oriented agent is essentially impacted in an inverse manner to an action oriented agent; its neuroticism factor is a function of the state orientation value and by extension weights the conation variable *self-relaxation*. The negative affect threshold to activate IM is increased while the positive affect threshold to activate IBC is lowered. Conversely to an action oriented agent, the agent fixates on negative outcomes. Unlike an action oriented agent, a state oriented agent derives greater satisfaction from external rewards, thus the reward any agent would receive is increased by a constant factor.

5) *Affective Regulation*: In PSI theory affect facilitates the shift between subsystem activation and each of the PSI subsystems map to one or more GR phases or motivational constructs. EM is represented by the motivational constructs while IM maps onto goals, and the formulation and planning phases (for complex plans). ORS is implemented as the situational evaluations, and IBC maps onto the execution and planning phases (for heuristic plans). Thus, to transition to the formulation or planning phase the agent must have a negative affective level, while to transition to the execution phase the agent must have a positive affective level; the management phase is active throughout the GR process and does not have affective requirements.

Alfred can execute two non-operative (non-op) actions, *beach stroll* and *meditation* which don't contribute as operators to any goal achievement directly but are used to employ the affect regulation techniques. By taking a beach stroll, Alfred is able to call the conation variable self-motivation, if it is accessible, and as a result increase its affective level. Similarly, by meditating, Alfred is able to call the conation variable self-relaxation, if it can be accessed, and as a result decrease its

Algorithm 2 Motivated GR Process

Input: Adopted goal set size, goal memory size, n GR cycle count, output file

Output: GR cycle output for m goals

```
for  $n$  goal reasoning cycles do
    rank goal library according to motivations;
    invoke baseline GR process with inputs;
    cache result := goal utility, outcome, name, plan, updated
    world state;
    update Alfred's internal state;
    update motive node with goal pursuit outcome;
end for
return goal reasoning process output for  $m$  goals
```

affective level. During goal pursuit, Alfred may not have the appropriate affective level to move to the subsequent phase, and may need to pause its goal pursuit activities with one of these non-op activities to manage its affective level.

TABLE I: Alfred's motivational profile for the *Cast Away* testing domain. The complete list of identity profile abbreviations can be found in table 5

Alfred's Identity Profile	
role_set:	functional robot, housekeeper
personality:	$C = 0.9, A = 0.9, N = 0.3, O = 0.7, E = 0.5$
belief_state:	boss(client, true), leisure_time(Alfred, 0)
value_system:	$sel = 0, sti = -0.75, hed = -1.0, ach = 1.0, pow = 0.0, uni = 0.5, sec = -0.25, con = 0.75, tra = -0.5, ben = 0.25$
need_satisfaction:	$aut = 0.25, ach = 0.25, pow = 0.1, aff = 0.25$
need_domain:	$family_home = 0.15, lab = 0.25$
drive_set:	self*, maintain order, control resources, discovery
action_orientation:	0.77
state_orientation:	0.14
volitional_competence:	0.23

V. EXPERIMENTATION

To test our adaptive self-concordant goal selection technique within the la VIDA system we performed an ablation test between the baseline and motivated GR processes. The agent, Alfred begins in a family home environment working as a home help bot for a client family, from which it is later separated during a sea voyage. In our test domain the Cast Away, we model this scenario by allowing the baseline GR process to goal reason over an assigned goal set **clean area**, **maintain self**, and **replenish energy**. After arriving on the deserted island Cast Away Isla, Alfred no longer has an assigned set of goals and can reason over the hypothetical "entire possible goal space" which consists of: **clean area**, **maintain self**, **replenish energy**, **cook**, **farm**, and **explore**. We run our ablation experiment using two objects *bgrp* of type GRProcess and *mgrp* of type MotivatedGRAgent. The *bgrp* object represents Alfred fulfilling its assigned set of goals in the family home environment under client supervision. While, the *mgrp* object represents Alfred selecting its own

goals from an arbitrary goal set while on Cast Away Isla. The goal_case_library is detailed in table IV, while the family home and Cast Away Isla initial world states are as shown in tables VI and V.

Except for the initial world states and goal libraries, the experimental configuration for *bgrp* and *mgrp* are identical. We have chosen to run the experiment, i.e., call goal_reasoning_process k times, to investigate how the system behaves on average for small numbers of goals. The max number of goals Alfred pursues per goal_reasoning_process call is determined by the gr_cycle_count variable. The actual experiment was run for $k = 100$, gr_cycle_cnt = 7, and adopted goal set size and goal_memory size equal to 2. When executed in a loop Alfred pursues 1 goal up to gr_cycle_count goals in groups which we refer to as goal pursuit sessions. The monitor and probabilistic flags were set to true to enable goal_reasoning_process to be executed from within the monitor strategy and the outcome of operator execution to be stochastic. The monitor strategy was then invoked within the enclosing k loop.

To evaluate its performance we extracted the accumulated goal utility for each goal pursuit session and summed it according to the goal count for k runs. Once all runs had been executed, each accumulated utility sum was divided by k , i.e., 100 to arrive at the average accumulated goal utility for 1, 2, ..., 7 goals as in figure 2. Although our results are far from conclusive, they are encouraging and suggest that the motivated GRProcess performs better overall than the baseline GRProcess. These results could be due to our starting values, e.g., goal utilities, and a next step would be to do more experiments with multiple randomly generated sets of values for the agent identity profile and goal library.




Fig. 2: Results of an ablation experiment comparing the accumulated utility in the family home domain for the baseline GR process and the cast away domain for the motivated GR process.

A. An analysis of Alfred's internal state

During the motivated goal reasoning process changes in Alfred's internal state were captured. Internal state changes are necessary for the motivation process and play an important role in motivation system activity according to PSI theory.

Alfred’s well-being as it changes over time is tracked by `accrued_wellbeing`. At the start of a motivated GR process, Alfred’s well-being is initialized to 0, but it is changed when Alfred gains well-being from achieving a goal, has positive affect, and additionally based on the successful or unsuccessful outcome of goal pursuit. According to the SCM, self-concordant goals and activities lead to an increase in well-being and increased effort and goals that are not self-concordant lead to a decrease in well-being and effort. Because `accrued_wellbeing` holds Alfred’s accumulated well-being and Alfred is trying to select goals that meet at least some of its need satisfaction requirements, we would expect well-being to trend upward which it does, this also implies that Alfred is experiencing sustained effort, otherwise no goals would be successfully achieved and well-being would trend downward over time.

According to PSI theory, positive and negative affect both play an important role in shifting between motivation subsystems; given our approach of modeling this behavior, Alfred will take actions to manage its affect when the threshold isn’t met to shift between GR phases. We observe that Alfred’s affective state oscillates, which is exactly what we expect; if Alfred had steady positive or negative affect, it would not be able to access each of the GR phases necessary to realize goal pursuit.

Frustration is defined as the ratio of failed subgoals to total subgoals pursued; according to the graph in figure 3 there was a peak in frustration at number of goals $n = 2$, but after which the graph generally trended down until $n = 7$ when there was a slight uptick. This is a positive outcome because it indicates that Alfred generally experienced sustained effort and failed less subgoals than it achieved. Reward has a similar impact on Alfred to well-being, except that it arises from external and introjected motivation; a reward is garnered from meeting external expectations. The importance of reward to and impact on the agent depends on its motivational configuration. Like well-being, reward is accumulated and thus could only be a monotonically increasing function, as the system currently has no negative rewards; the graph in Fig. 3 depicts a consistent increase in Alfred’s rewards.

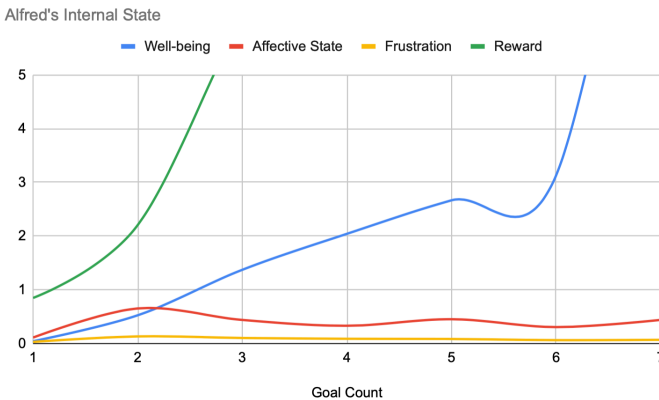


Fig. 3: Alfred’s internal state as experienced in the baseline vs. motivated ablation experiment.

VI. CONCLUSION

This work explored a solution to the problem of an agent self-selecting a new set of objectives after a previously assigned goal set has been invalidated. Because communication with a controlling entity may be limited, the agent will need some other means to guide the goal pursuit process. Our solution implements a goal reasoning formalism as the foundation of a motivated goal reasoning system modeling the human self-system elements roles, beliefs, personality, values, and autobiographical memories with two motivation components called an identity profile and motive network. These elements interact to allow the agent to pursue useful goals while also having behavior that aligns with its own motivation system.

The agent’s human-inspired motivation system allows us to leverage the theory of goal self-concordance for human self-regulation to select aligned goals while managing agent effort and well-being. The SCM posits that the selection of self-concordant goals in self-concordant contexts leads to more sustained effort and well-being, thus we expected that our experiment results would indicate higher average utility when the agent is augmented by a motivation system. In addition it was expected that high utility goals, which harmonized with the agent’s motivations and abilities would lead to less goal failures. With the preliminary testing conducted this is what was observed, however more testing is needed to draw any strong conclusions.

This work makes a number of important contributions which we briefly list and describe below.

- An implementation of the GR process developed by [Roberts et al., 2016], in which we use our own custom combination of refinement strategies, goal and task decomposition methods, operators, and algorithm for executing a goal reasoning process.
- A model and concrete implementation of our novel identity profile and motive networks to characterize a simplified version of the human self-system.
- A set of conation strategies that act as a meta-reasoning process and guides the agent’s goal reasoning pursuit according to PSI theory.
- An adaptive self-concordant goal selection technique based on SCM theory
- A hybrid motivated goal reasoning process that augments the baseline process with the agent’s motivation system.

The key roles of the motivation system as the system is currently implemented is in the ranking of goals and managing of the agent’s internal state. When a goal class is pursued for the first time, the agent determines the alignment with the goal and itself by calculating the goal self-concordance of the goal. But, if the goal has already been pursued the agent has an advantage because it can reduce the uncertainty in the goal selection process. Our method handles the uncertainty in establishing new goals when the agent’s environment provides limited guidance, by having a mechanism to help determine behavior in these types of situations, and it reduces uncertainty

in how the agent may approach a category of unfamiliar situations.

Because this work implements the first iteration of our system as a proof of concept, we have made a number of design choices that are not scalable. The plan library, goal library, and common sense repository are currently developed by hand requiring an intimate knowledge of how the system works to create the parts correctly. Other caveats include a need for more extensive testing, currently we conduct an ablation test comparing an implementation of a GR formalism with a modified version that integrates our motivational constructs and processes; the sole metric for these tests is utility, because motivation metrics are not available for the baseline GR process. While our results can hint at the usefulness of la VIDA as it is implemented, more testing and usage is needed to uncover the extent of its shortcomings.

In future research we would like to work on making the plan library, goal library, and common sense repository more useful, automated, and designer friendly. Because the GoalNode data structure is a GTN we would like to investigate replacing our plan library with an implementation of GTPyhop [Nau et al., 2021] or similar planner. And finally, to test our agent more thoroughly and discover additional research directions, we would like to deploy it in a rich virtual environment where there are a combination of artificial and human agents executing purposeful behavior.

REFERENCES

- Ursula Addison. The source of desire: Personal identity as a drive for agent behavior. In *AIC*, volume 3400 of *CEUR Workshop Proceedings*, pages 147–155. CEUR-WS.org, 2022.
- Ursula Addison. Human-inspired goal reasoning implementations: A survey. *Cognitive Systems Research*, page 101181, 2023. ISSN 1389-0417. doi: <https://doi.org/10.1016/j.cogsys.2023.101181>. URL <https://www.sciencedirect.com/science/article/pii/S1389041723001158>.
- David W. Aha. Goal reasoning: Foundations, emerging applications, and prospects. *AI Magazine*, 39(2):3–24, Jul. 2018.
- Ron Alford, Vikas Shivashankar, Mark Roberts, Jeremy Frank, and David W. Aha. Hierarchical planning: Relating task and goal decomposition with task sharing. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI’16*, page 3022–3028. AAAI Press, 2016.
- John Anderson. *ACT: A Simple Theory of Complex Cognition*, pages 49–70. 09 2002. ISBN 9780262281744. doi: 10.7551/mitpress/1888.003.0006.
- J. Boyd, G.T. Hammond, and Air University (U.S.). Press. *A Discourse on Winning and Losing*. Air University Press, Curtis E. LeMay Center for Doctrine Development and Education, 2018. ISBN 9781585662791. URL <https://books.google.com/books?id=B-aVuQEACAAJ>.
- Alexandra Coman and Hector Muñoz-Avila. Motivation discrepancies for rebel agents: Towards a framework for case-based goal-driven autonomy for character believability. *ICCBR-14 Workshop on Case-based Agent*, 2014.
- T Cox, Michael. A model of planning, action, and interpretation with goal reasoning. *Advances in Cognitive Systems*, 5:57 – 76, September 2017.
- Bart Gajderowicz, Mark Fox, and Michael Grüninger. Requirements for emulating homeless client behaviour. 02 2017.
- Bart Gajderowicz, Mark S. Fox, and Michael Grüninger. The role of goal ranking and mood-based utility in dynamic replanning strategies. *Advances in Cognitive Systems*, 9: 211–230, Sept 2018.
- Nick Hawes. A survey of motivation frameworks for intelligent systems. *Artificial Intelligence*, 175(5-6):1020 – 1036, 2011.
- Jutta Heckhausen and Heinz Heckhausen. *Motivation and action*. Cambridge University Press, 2008.
- Pierce Howard and Jane Howard. The big five quickstart: an introduction to the five-factor model of personality for human resource professionals. 1995.
- Matthew Evans Klenk. Goal-driven autonomy in planning and acting. 2010.
- Julius Kuhl. Chapter 5 - a functional-design approach to motivation and self-regulation: The dynamics of personality systems interactions. In *Handbook of Self-Regulation*, pages 111–169. Elsevier Inc, 2000. ISBN 0121098907.
- Julius Kuhl and Nicola Baumann. Chapter 27 - personality systems interactions (psi theory): Toward a dynamic integration of personality theories. In John F. Rauth-

- mann, editor, *The Handbook of Personality Dynamics and Processes*, pages 709–730. Academic Press, 2021. doi: <https://doi.org/10.1016/B978-0-12-813995-0.00027-3>.
- John E. Laird. *The Soar Cognitive Architecture*. The MIT Press, 2012. ISBN 0262122960.
- Traci Mann, Denise de Ridder, and Kentaro Fujita. Self-regulation of health behavior: social psychological approaches to goal setting and goal striving. *Health Psychology*, 32(5):487–498, 2013.
- Marina Milyavskaya, Daniel Nadolny, and Richard Koestner. Where do self-concordant goals come from? the role of domain-specific psychological need satisfaction. *Personality and Social Psychology Bulletin*, 40(6):700–711, 2014. doi: [10.1177/0146167214524445](https://doi.org/10.1177/0146167214524445). URL <https://doi.org/10.1177/0146167214524445>. PMID: 24625657.
- Hector Muñoz-Avila. Adaptive goal driven autonomy. *Case-Based Reasoning Research and Development. ICCBR*, 11156:3–12, 2018.
- Dana Nau, Sunandita Patra, Mark Roberts, Yash Bansod, and Ruoxi Li. Gtpphop: A hierarchical goal+task planner implemented in python. *ICAPS Workshop on Hierarchical Planning (HPlan)*, August 2021.
- Mark Roberts, Vikas Shivashankar, Ron Alford, Michael Leece, Shubham Gupta, and David W. Aha. Actorsim: A toolkit for studying goal reasoning, planning, and acting. *Advances in Cognitive Systems*, June 2016.
- Paul S. Rosenbloom, Abram Demski, and Volkan Ustun. The sigma cognitive architecture and system: Towards functionally elegant grand unification. *Journal of Artificial General Intelligence*, 7(1):1–103, 2017. doi: [10.1515/jagi-2016-0001](https://doi.org/10.1515/jagi-2016-0001). URL <https://doi.org/10.1515/jagi-2016-0001>.
- Alexei V. Samsonovich. Emotional biologically inspired cognitive architecture. In *Biologically Inspired Cognitive Architectures*, 2013.
- R.C. Schank and R.P. Abelson. *Scripts, Plans, Goals, and Understanding: An Inquiry Into Human Knowledge Structures*. Artificial intelligence series. L. Erlbaum Associates, 1977. ISBN 9780470990339. URL <https://books.google.com/books?id=YZ99AAAAMAAJ>.
- Shalom H. Schwartz. An overview of the schwartz theory of basic values. *Online Readings in Psychology and Culture*, 2:11, 2012.
- Kennon M. Sheldon. Becoming oneself: The central role of self-concordant goal selection. *Personality and Social Psychology Review*, 18(4):349–365, 2014. doi: [10.1177/1088868314538549](https://doi.org/10.1177/1088868314538549). URL <https://doi.org/10.1177/1088868314538549>.
- Kennon M. Sheldon and Andrew J. Elliot. Goal striving, need satisfaction, and longitudinal well-being: The self-concordance model. *Journal of Personality and Social Psychology*, 76(3):482–497, 1999.
- Ron Sun. Motivational representations within a computational cognitive architecture. *Cogn. Comput.*, 1(1):91–103, 2009.
- Daniel Swoboda, Till Hofmann, Tarik Viehmann, and Gerhard Lakemeyer. Towards using promises for multi-agent cooperation in goal reasoning. *2022 Workshop on Planning and Robotics*, 2022.
- Swaroop Vattam, E. Klenk, Matthew, Matthew Molineaux, and David W. Aha. Breadth of approaches to goal reasoning: A research survey.
- Xinmiao Yu, Riccardo Morri, and Fernanda M. Elliott. Eda, an empathy-driven computational architecture. *Proceedings of the Ninth Goal Reasoning Workshop*, (9), Nov 2021.

APPENDIX A ABLATION STUDY INITIAL STATES

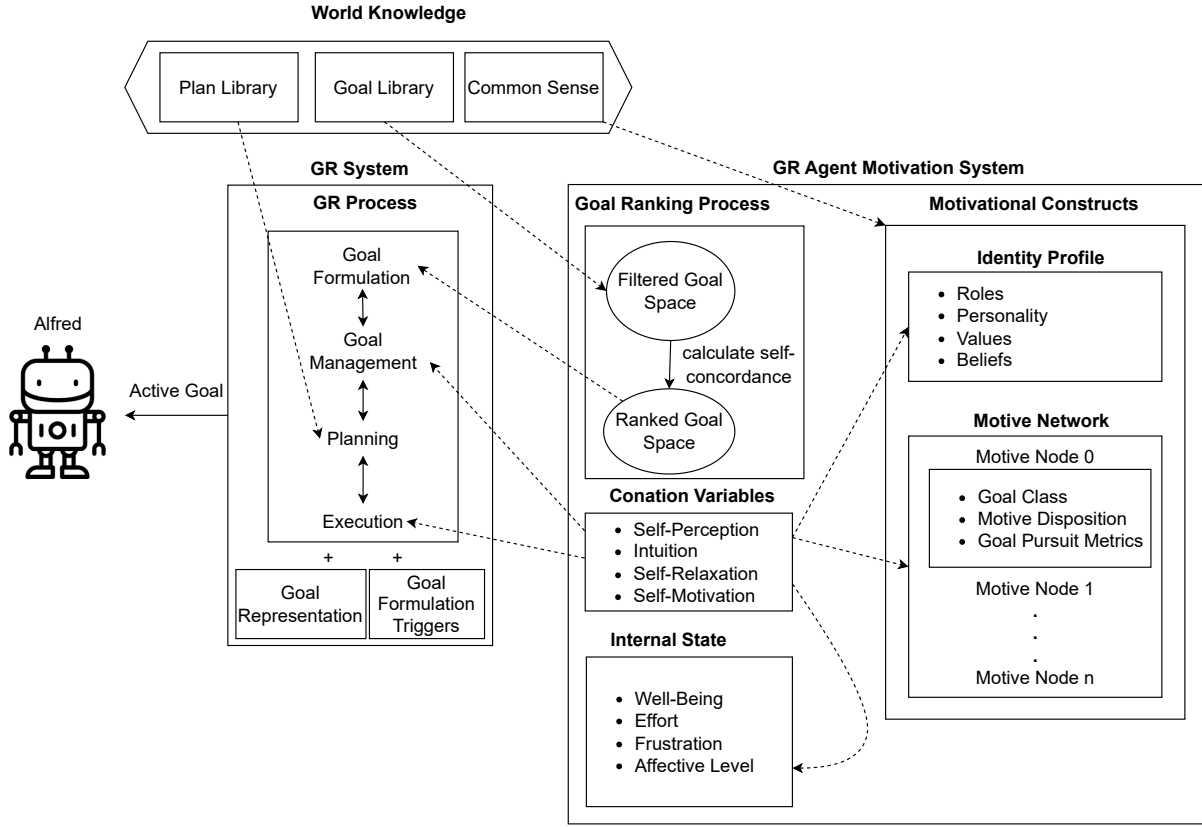


Fig. 4: Displayed is an overview of the la VIDA system, where GR process is the goal reasoning mechanism, the GR agent motivation system holds the motivation data structures and implements the conation variables, and the world knowledge provides some semantic information for motivation elements, as well as possible goals and their plan recipes.

PlanLibrary Function Semantics	
m_move	move an object to a destination (put, walk)
m_toggle_power	turn a machine on or off (toggle_power_switch)
m_machine_use	operate and get the result of using a machine (recharge_battery)
m_clean	clean entity (uv_flash, wipe, vacuum, landscape)
m_update_software	put software in functional state (install_software)
m_hardware_repair	put hardware in functional state (humanoid_repair)
m_prepare_meal	prep and cook raw food (indoor_food_prep, outdoor_food_prep)
m_farm_crop	care for plant (plant_care)

TABLE II: The purpose of each task method and the set of operators it can return.

Cast Away Domain Goal Space	
Replenish energy	refills Alfred's current energy level to full
Maintain self	performs basic upkeep to keep Alfred in functional health
Clean area	cleans a region
Cook	preps raw food items and cooks them
Farm	cares for a plant to take it from seed to fruit status
Explore	visits a previously unexplored region

TABLE III: Possible goals Alfred can pursue in the family home or cast away island environments.

Personality	Conscientiousness	Agreeableness	Neuroticism	Openness	Extraversion
Value System	Self-direction	Stimulation	Hedonism	Power	Security
	Achievement	Conformity	Tradition	Benevolence	Universalism
Need Satisfaction	Autonomy	Affiliation	Power	Achievement	

Fig. 5: Motivational entries for an identity profile. Bolded letters indicate element abbreviations.

TABLE IV: Goal case values for goal classes in the Cast Away domain, all numerical values were randomly generated initially and reused there after. The entries ns_ach and ns_pow are need satisfaction components, while vs_ach and vs_pow are elements of the value system; the complete list of abbreviations can be found in table 5.

Cast Away Goal Case Library						
Goal class	cook	farm	explore	clean area	maintain self	replenish energy
ns_ach	0.01	0.12	0.14	0.03	0.01	0.05
aff	0.21	0.17	0.25	0.21	0.09	0.23
aut	0.04	0.12	0.11	0.2	0.2	0.1
ns_pow	0.23	0.06	0.21	0.18	0.04	0.22
difficulty	1	2	3	2	4	5
urgency	3.0	2.7	2.55	3.0	0.05	1.0
utility	17	16	6	10	13	4
need domain	family home	outdoors	outdoors	family home	lab	family home
C	0.46	0.63	0.83	0.29	0.65	0.43
A	0.84	0.85	0.89	0.13	0.23	0.0
N	0.19	0.56	0.5	0.18	0.23	0.98
O	0.67	0.35	0.81	0.83	0.29	0.76
E	0.63	0.77	0.12	0.05	0.03	0.18
vs_ach	0.78	-0.37	0.86	0.14	-0.33	0.06
con	0.99	0.38	0.01	0.56	-0.19	-0.69
sel	-0.45	0.54	0.11	0.91	-0.18	0.64
vs_pow	-0.79	-0.28	1.0	-0.47	0.56	0.94
sec	-0.08	0.58	0.7	-0.05	0.47	0.89
ben	0.38	-0.76	-0.99	0.79	0.45	-0.25
uni	-0.1	0.58	0.49	-0.75	-0.41	0.96
sti	-0.22	0.7	0.32	-0.94	-0.35	-0.97
tra	-0.97	-0.59	0.65	0.38	-0.7	0.5
hed	-0.88	-0.97	-0.74	-0.97	0.54	0.52
drive	control resources	control resources	discovery	maintain order	self*	self*

