Character-Oriented Narrative Goal Reasoning in Autonomous Actors

Alexei V. Samsonovich ASAMSONO@GMU.EDU
Krasnow Institute, George Mason University, University Drive, MS 2A1, Fairfax, VA 22030 USA

David W. Aha DAVID.AHA@NRL.NAVY.MIL
Navy Center for Applied Research in Artificial Intelligence, Naval Research Laboratory (Code 5514), Washington, DC 20375 USA

Abstract

The cognitive autonomy capability understood in a certain sense is vital for actors working as teams in unpredictable environments with limited communications. Such actors need to be able to reason about their goals and decide which to pursue. Therefore, methods of autonomous goal reasoning (GR) would be useful in such cases. However, most previously studied models of GR are limited in their abilities to reason in situations involving multiple characters, viewing the latter from a third-person perspective. In this context, the connection between GR and narrative reasoning is beginning to be explored. Recently, a theoretical framework for narrative team planning (NTP) has been developed and used to generate character goals (not team goals). This approach, however, relies on a global team planning, which may not be suitable for a team operating with limited communications and without a central command. Accepting NTP as a baseline, here we argue that an appropriate introduction of characters and narratives into a GR framework can (in some environments) be beneficial, both conceptually and in terms of performance measures. We study this claim analytically using a Character Reasoner model that we formulate, and apply it to several example scenarios. This model separates characters from actors and employs a hierarchical network for narrative and goal selection. Expected benefits for teams of autonomous actors include more efficient and more robust goal reasoning abilities.

Keywords: Goal reasoning; autonomy; narrative intelligence; character arc; metacognition; distributed multi-agent reasoning

1. Introduction

Goal reasoning (GR) refers to the ability of a cognitive system to deliberate on, generate, and select its own goals in unforeseen situations (Vattam et al., 2013; Klenk et al., 2013; Roberts et al., 2014). Unlike most research on intelligent agents, GR is inspired by observations of (highly autonomous) human cognitive behavior. The high level of cognitive autonomy provided by GR can be vital for actors working in teams and/or in unpredictable environments. As an illustration, we consider a simple scenario (please see Sections 3 and 5.1 for further details of this example).

Suppose a team of two combat pilots are engaging an enemy in a 2-versus-2 beyond-visual-range air combat scenario. At the same moment, they receive information from their commanding officer about a developing situation at a distant location. They infer that the immediate help of at least one air vehicle is vital to that situation, is more important than their current mission, and requires a vehicle that is fully armed. The team needs to reassess, and possibly formulate, goal(s) to pursue. On the one hand, they cannot simply abandon the ongoing engagement: disengaging may result in the destruction of friendly assets. Therefore, one option is to pursue the previous goal until it is achieved, then switch to the new goal. However, in this case none of the team members will be guaranteed to have a full load of offensive assets, so their help at another location could be useless. An alternative choice is to arrange for one of the vehicles to escape before firing any of its missiles. For this to happen, the pilot of the remaining aircraft must keep both enemies engaged, allowing the partner to escape (while still fully loaded). Once the two goals are formulated, the associated risks and values can be estimated, and the selection among them can be made. This task requires more than traditional planning: for example, the team needs to select a top-level goal at the start of their engagement. Moreover, the need for team-level GR may repeatedly be required in a dynamically changing scenario, as we exemplify below. Existing GR agents aren't designed for these kinds of open-ended team decision-making scenarios.

While GR has only recently been an active research topic, many branches of AI research relate to it. Examples include topics like intelligent agents and cognitive architectures (Gray, 2007), motivated agents, narrative planning, intent and plan recognition, self-regulated learning (Zimmerman, 2002), and many more. Because it is not our present goal to review all these areas, we contrast our approach to only one related topic, namely *narrative planning* (Riedl et al., 2008; O'Neil & Riedl, 2014; Riedl & Young, 2010, 2014; Young et al., 2013) applied in the context of a team mission. We refer to it as *narrative team planning* (NTP) and take it as a baseline.

The essence of NTP can be outlined as follows. First, the narrative that determines future team actions is generated by a planner, given one desired outcome as the goal (some systems allow for goal replacement in the case of a plan generation failure; see (Riedl et al., 2008)). Second, the actions of each character must be believable (i.e., consistent with individual character goals and motives, which are not necessarily consistent with each other or with the team's goal). This constraint allows the actors to execute plans locally within the team. Here "locally" means focusing on their individual tasks and tasks of selected peers rather than on the entire team mission. Third, plans are generated from a "global" team perspective. The global narrative is created by eliminating unmotivated commitments (e.g., by assigning certain individual goals to characters, which transforms NTP into a GR process). Finally, the set of NTP characters is known a priori and is fixed, because the term "character" is synonymous to the term "actor" in NTP.

The approach that we introduce here as an alternative to NTP can be characterized as "character-oriented narrative goal reasoning", or, for short, "character reasoning" (CR). A character in this case is an abstraction, and is not identical to an actor. We define a *Character Reasoner* as an autonomous, embodied intelligent agent (an actor) capable of (a) formulating believable characters applicable to the developing situation, that help with achieving team's goals, and (b) performing roles of selected formulated characters, using narrative GR and planning in cooperation with the team. CR relies on the concepts of a narrative, a character, a character arc (i.e., a plot or a storyline describing the evolution of a character and its goals in a story), and related concepts, which are useful in many (e.g., military) domains (Finlayson & Corman, 2013) and are also parts of the NTP formalism (Riedl & Young, 2010; Ware & Young, 2014). So, what is essentially new in CR, besides separation of characters from actors?

To answer this question, we shall return to our example, assuming now that there are $N \geq 2$ vehicles in each team, M vehicles need to escape without firing any of their missiles, and $1 \leq M < N$. As explained above, for the air combat team to select a new goal using NTP, the future actions of all team members (actors) need to be considered and optimized as a whole. This could make the task computationally demanding and practically intractable in the case of a large team size. Moreover, the necessary information about all partners may not be available locally to team members engaged in NTP. Therefore, an important question is whether and how the global team-GR task can be effectively decomposed into individual, local GR tasks. This is exactly what the present work intends to address (Section 3 explains intuitively how, using the selected example).

The paper is organized as follows. In Section 3 we show how to perform the decompositions discussed above, but only after we provide a minimal background on the topic in Section 2. We also explain in Section 3 why CR can be more efficient than NTP, using the selected example at an intuitive level. We then formalize CR in Section 4, casting it as a form of narrative GR, and consider its likely benefits based on a more detailed analysis of other examples in Section 5. We finally discuss implications and summarize the main points in Sections 6 and 7.

2. Background and related works

In narratology, a *narrative* is defined to have two related components (Bal, 1998; Riedl & Young, 2010; Schmid, 2010): the *fabula*, which is a partially ordered set of causally, logically or temporally related events that form a consistent story, and the *sjuzet*, which is a (possibly incomplete) linear presentation of this story as viewed from a storyteller perspective[1]. Narrative techniques and their application have received recent attention in AI research, yet not sufficiently for team-level GR and planning and military applications (see, however, Young et al., 2013; Finlayson & Corman, 2013). At the same time, there is no general consensus on a formal definition of narratives, and how they should be generated (Kapadia et al., 2015). Within the popular formalism developed by Young's research group (Riedl et al., 2008; Riedl & Young, 2010, 2014; Young et al., 2013, Ware & Young, 2014), a fabula is defined as a sound plan with the additional requirement of character believability, enforced through the consistency of character goals, intentions and actions. Accordingly, within this approach, narratives are generated by planning algorithms. However, this use of planning imposes severe limitations on the outcome, many of which are identified by the authors themselves (Riedl & Young, 2010).

Outside this relatively narrow understanding of a narrative, no precise criteria are defined to distinguish narrative and non-narrative planning or GR. For example, John McCarthy defined a narrative more loosely than a plan: as a temporal partially-ordered collection of related situations and events, without requiring their consistency (McCarthy, 1994; McCarthy & Costello, 1998). In the situation calculus, narratives are first-order objects. From this point of view, today virtually any symbolic cognitive architecture can be regarded as a narrative-based modeling framework. McCarthy (1994; 1998) also discussed the concept of a *proper narrative,* which is a narrative without anomalies. Among other formal approaches in narratology, Abell (1987; 1993) represents a narrative as a graph, the nodes of which represent actions performed by actors that change states

---

[1] This dichotomy originates from the Russian mechanistic formalism of literary criticism developed early in the 20th century. Other terms are also used (e.g., "story", "discourse", or "plot").

of the world. Edges of the graph are directed and represent causal or dependency relations, indicating, for example, that one action is a prerequisite for another. Some modern authors go further and require that a narrative should be understood more narrowly than a plan constrained by the believability of characters. Additional requirements include "narrativity", character interaction, a conflict or suspense, struggle, characters changing their goals and values, and more (Huhn, 2013; Simon-Shoshan 2013; Haven, 2007). Finlayson and Corman (2013) refer to this kind of a narrative as a Level-II narrative.

In this paper, we distinguish between CR and non-CR approaches (we define CR formally in Section 3.2) using the notion of a character, which we contrast with the notion of an actor. For Haven (2007), characters are abstractions that are central in understanding a narrative, or story. Haven defines a story as a detailed, character-based narration of a character's struggle to overcome obstacles and reach an important goal. A character, according to Haven (2014), is an abstraction defined by five elements that can be interpreted (in our words) as follows:

1. drives, values and motives (the *core*);
2. personality type and features;
3. autobiographical memory and general knowledge;
4. behavioral activity and capabilities; and
5. subjective viewpoint, appearance, and self-image.

A character in a narrative is a perspective and a viewpoint that allows us "to see who is doing the action and to gauge relevancy…", "to interpret emotional state, beliefs, attitudes… to create meaning and relevance". A character's behavior is necessarily intentional. Character *intent* is composed of two components: the *goal* (the outcome that is being pursued by the character) and the *motive* (i.e., the reason why this goal is important for the character) (Haven, 2007).

Thus understood, characters are distinguished from actors (intelligent agents that are given entities), as well as from cognitive systems and from objects in the environment. A *character* is understood in this Section and below as an abstraction in the form of a virtual subject (an ego, a self, a persona), defined by its subjective perspective, together with a system of values, motives (given by top-level guidance and/or bottom-level drives), beliefs (autobiographical and general), and capabilities. A character is associated with a particular *arc* in a narrative, and can be performed by an actor.

Although both the character and the actor could be virtual agents, the notion of an actor is not redundant in this context. An actor is a fixed entity for a given specific scenario, while characters can be dynamically created, modified and deleted, as they exist in the "minds" of actors. An actor may have a suit of characters and make selections among them depending on the situation. In this context, "character" is a synonym of "role". For actors, characters replace goals and commitments, while being richer and more powerful constructs. Playing a certain character means more than pursuing a certain goal. Typically it involves an evolution of the character's goal(s). This notion of a character is further illustrated using a set of example scenarios in Section 5.

The primary distinction to be made between the present work and related works is that this work proposes to use characters as specifically defined narrative structures (distinct from goals and actors) to assist GR in complex scenarios (e.g., involving multiple actors and characters).

3. Local decomposition of team-GR using CR: An intuitive preamble

As the name suggests, character reasoning involves the concepts of a character and a character arc. As mentioned in Section 1, we distinguish characters from actors, which is not the case in NTP. A *character* in CR is an abstraction: it is a virtual rational agent with its own goals, motives, senses, affordances, knowledge, and recent history. In general, there is no 1-1 correspondence between actors and characters in CR. One actor can play multiple characters in a sequence, and in certain cases in parallel. The purpose of introducing characters in CR is to decompose the global team-level GR into local, single-character GR. This decomposition involves two steps: (1) separation of the global GR task into two stages (in the first of which characters and prototype character arcs are selected that will form the narrative), and (2) decomposition of the second stage (which maps characters to actors and further specifies character arc details) into individual-character GR tasks (Figure 1). Therefore, in this character-oriented version of GR, the notion of a character is central and includes the notion of a goal.

To illustrate how CR works, we extend our example air combat scenario. Two types of characters can be identified in our example: one, a “deserter”, who escapes the fight without deploying any assets, and another, a “hero”, who keeps the enemy engaged. These characters are believable (they act according to their goals and motives), yet their goals, motives, and strategies differ. Once defined as a part of the narrative, these two character types essentially determine the team’s strategy (we assume that this decision is made independently by both pilots at the beginning). Yet, the character mapping to actors may not be decided immediately, is ambiguous, and may change over time. If all combat vehicles are identical and the pilots have similar expertise, then the choice should be determined by their relative positions. In the absence of a central command, team-level GR is performed independently by each actor, who must decide which character to play. The choice of each actor should be consistent with the team, and may need to be corrected based on observations of the partner’s behavior or communications. For example, if you are an actor in the team, then upon noticing that your partner’s behavior is inconsistent with your character choice, you can adjust your choice or communicate with your partner (i.e., to convince them to change their character choice). Once the mapping of characters to actors is decided by the team, each actor assumes the goal, motives, etc. of the selected character, and starts planning and acting based on this choice, possibly while keeping only marginal awareness of the rest of the team. In other words, the task becomes locally decomposed.

Both stages of CR in our example are executed locally. Stage-1 CR is performed by each actor independently in parallel, producing the same result (deciding that the two character types will determine the narrative: suppose that actors produce the same set of character types). Stage 2 involves local communications (Figure 1). This stage may recur during actor-level plan execution. Indeed, imagine that at a certain point swapping the characters becomes advantageous for the team (e.g., if both enemies follow the deserter). Character swapping can be performed using a global team-level GR. However, it can also be done using local communications, if we add the “swapping affordance” as a special privileged action to the deserter’s repertoire (instantly switch positions with the partner). When feasible, this action should also transfer the character’s current plan and recent memory from one actor to another. This will allow each character to preserve continuity, while reasoning locally (we consider the jump as a local action). In either case, the task remains decomposed at all times.

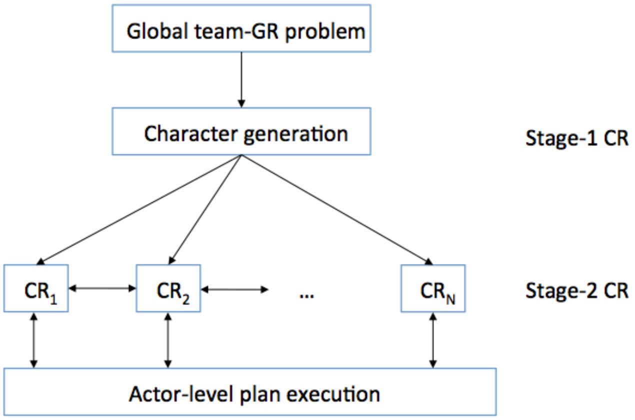

Figure 1. A general decomposition scheme of the team-level GR problem using character reasoning for the *N*-aircraft example. Each CR$_i$ is an individual-character GR process, including the mapping of the character to an actor. Horizontal arrows show local communications, possibly including the swapping of characters.

The practical difference between CR and NTP in this example may be small, given that the number N of aircraft is 2, but it could be noticeable when N is large. For example, suppose there must be M deserters and $N - M$ heroes, where $M \sim \frac{N}{2} \gg 1$. Then the number of choices

$$\binom{N}{M} = \frac{N!}{M!(N-M)!}$$

grows superexponentially with N, and could become too large to explore in real time using NTP. In contrast, if using CR all choices can be made locally in parallel. While there is some tradeoff (communications among team members are required), this does not necessarily assume all-to-all individual communications. For example, suppose that a nearly random mapping of character types to actors is accepted by the team initially, when each actor decides independently which character to play based on his local surrounding, and broadcasts his choice to the team. Then, further optimization of the mapping can be obtained using local character swapping. As a result, the team can be expected to produce nearly optimal behaviors for large N using CR. This illustrates a primary distinction of CR and NTP, and applies to other examples.

4. General conceptual basis and the CR formalism

In this section we define CR and its main building blocks formally. We start from a top structure that we call a *hierarchical narrative network* (HNN), related to the notion of a narrative network (Pentland & Feldman, 2007), which is defined as the tuple:

$$\text{HNN} = \langle S, E, C, \mathcal{A}, P \rangle, \qquad \text{(Eq. 1)}$$

where S is a set of nodes, E is a set of directed edges, C is a set of characters, $\mathcal{A}$ is a set of character arcs, and P is a set of performing actors. Now we will explain these elements intuitively. An HNN includes a graph with a set of nodes S and a set of directed edges E. Here nodes represent actual and possible states (see the definition of a state below in this section) and fragments (i.e., fragments of the graph), and edges represent causal and temporal relations among nodes, inducing a partial order. Possible states, relations and rules of dynamics are given by the *domain theory*: a knowledge base that is assumed given, yet is not explicitly included in (Eq. 1). The network is hierarchical, because some of its nodes represent fragments of the same network. Fragments can be collapsed into nodes, and nodes can be expanded into fragments, as necessary.

In addition to S and E, an HNN (Eq. 1) includes a set of characters $C=\{c_j\}$, a set of possible character arcs $\mathscr{A}=\{A_j\}$, and the set of performing actors $P = \{P_i\}$. Our intuitive notion of a character was introduced in Section 2. Technically, a character c is represented by a tuple

$$c = < p,\ m,\ A >, \qquad \text{(Eq. 2)}$$

in which the character is given by the perspective p, motives m, and the arc A. Now we shall explain these terms. The perspective p represents the current character's viewpoint, including senses of "now" and "here", "self" and "others" (own identity), as well as any specific features and capabilities, and other contextual variables determined by embodiment (i.e., the performing actor P_i). The set of character's motives m includes drives, values, and top-level guidance, that determine the selection of goals and intentions and usually do not change their nature within a character arc. A character arch was defined intuitively in Section 2. Formally, the character arc A is a set of character's attitudes $\{a_i\}$, such as beliefs, goals, intentions, memories, percepts and affordances, taken as functions of time. In general, a character c's attitudes are formed from states by attributing them to the character together with a certain modifier, e.g.:

```
c.does(s0), c.intends(s1), c.ignored(s2),
c.achieved(s3), c.saw(s4), c.committed(s5).
```

A character in an HNN is not committed to a particular arc or actor, and therefore may not have a unique perspective. However, each character c in a given narrative is committed to an arc A and to an actor p (the embodiment of c). In the spirit of Abell's formalism (2009; 2011), we define a *fabula* as any part of an HNN written as the tuple (Eq. 1) that is internally closed (i.e., all intentions and actions are *motivated* and placed into arcs), consistent, and includes exactly one arc per character. Here "motivated" means that character intentions can be explained by, or derived from character motives. Also, a mapping of performing actors to characters needs to be specified in the fabula. Then, we define the *sjuzet* to be the arc corresponding to the storyteller (usually the protagonist or the author; in our case of interest it is the character of the reasoning actor).

We define a state $s \in S$ to be a class of possible physical states of the world, such that a particular given fact applies to all those and only those physical states of the world. For example, a state can corresponds to a specified place and/or a moment or an interval in time, and/or can represent a particular object that is present there, or an event, a condition, a feature or property, etc., and any collection of them. States are therefore not necessarily mutually exclusive: for example, there may be more than one actual current state represented in the actor's working memory. Fractions and unions of states are themselves states. Therefore, states can be added and subtracted as sets. Controlled actions and processes are also states. Due to the hierarchical nature of HNNs, HNN fragments are also states, when represented by nodes. When a state is included in a narrative, it allows characters to form attitudes based on this state.

Finally, we define the term *Character Reasoner* as a system that implements the formalism described in Equations 1 and 2, and operates on an HNN, producing characters, character arcs, and a narrative. The architecture of a Character Reasoner is shown in Figure 2. In our case, the Character Reasoner is the author, the storyteller and the actor at the same time (e.g., the sjuzet is presented from the perspective of its character). The general CR procedure is outlined below (here we assume that the set of performing actors *P* and the domain theory *D* are given).

The CR procedure:

(1) Start with populating/updating HNN states and relations, using the available input and referring to the domain theory.
(2) Formulate relevant possible characters. Select characters that will determine a narrative.
(3) Generate possible expected character arcs for selected characters.
(4) Combine the arcs into a set of possible narratives (fabula).
(5) Evaluate and compare generated narratives, select the working narrative.
(6) Map characters in the working narrative to actors.
(7) Generate a sjuzet with the selected character(s) as the storyteller(s).
(8) Use the sjuzet to plan and execute the task at the actor level.

In a team, the outcomes of (2), (5), and (6) should be confirmed by partners (either explicitly, via communications, or implicitly, by observable behavior). Conflicts need to be resolved before plan execution.

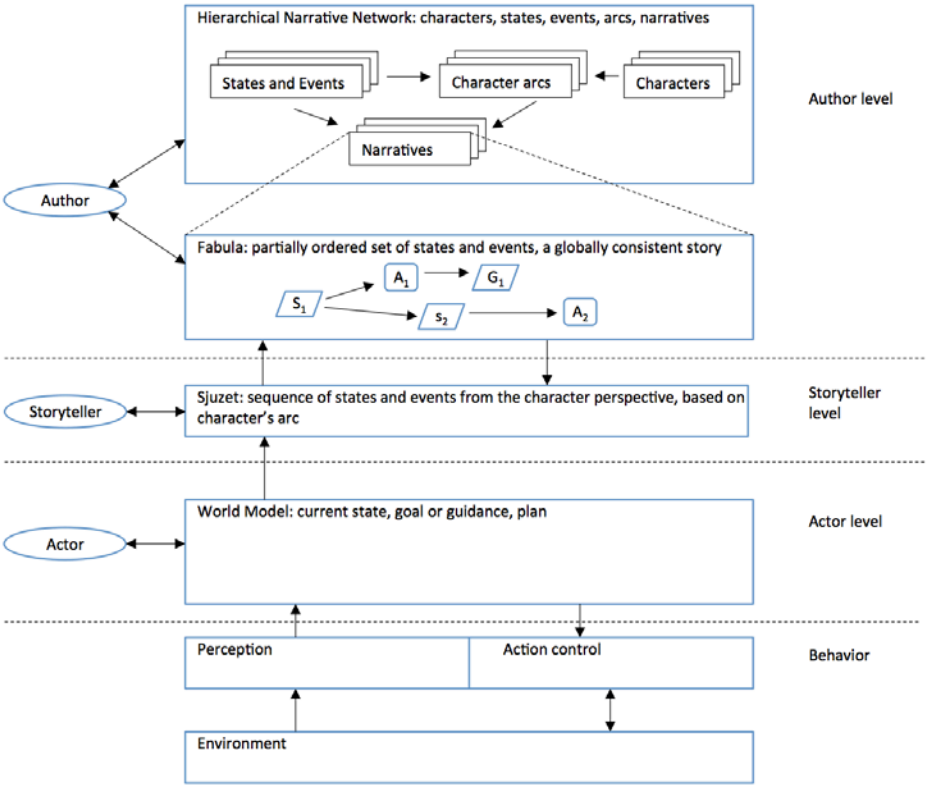

Figure 2. General architecture of a Character Reasoner that is also an actor.

5. Analysis of examples

To assess the scope and benefits of our formalism, we consider the following examples at a more detailed level, each formulated in a virtual environment, allowing for their future computer simulations that will be presented elsewhere. We start with the example already used above.

5.1 Air combat (one possible scenario)

Initial situation: the Blue team consisting of two fighter jets, Snake1 and Snake2, are patrolling a political border. The Red team (bandits), consisting of two enemy aircrafts, has just crossed the border. Each aircraft carries two long-range missiles (we assume that no other weapons are used). The bandits fly close to each other, while Snake1 and Snake2 fly at a distance from each other. Commanding officer informed Snake1 and Snake2 that two bandits crossed the border and need to be eliminated. The initial distance between the two teams is beyond the radar / missile range.

Original working narrative:
T1. Bandits search for Blue aircraft using their radars.
T2. Snake1 and Snake2 search for bandits using their radars.
T3. Snake2 detects Red radars and communicates to Snake1.
T4. Snake1 detects Red radars and communicates to Snake2.
T5. Snake1 and Snake2 triangulate Red coordinates.
T6. Bandits detect Blue radars.
T7. Snake2 snoozes the radar.
T8. Snake1 activates a radar jammer.
T9. Bandits activate jammers.
T10. Bandits turn to pursue Snake1 but cannot achieve a missile shot.
T11. Snake2 fires two long-range missiles at bandits.
T12. One of the two bandits is destroyed.
T13. Snake2 illuminates the remaining bandit with a radar, as a distraction.
T14. The bandit pursues Snake2.
T15. Snake1 takes a shot at the remaining bandit.
T16. The bandit is destroyed.

Working narrative altered at T10, no character-actor separation:
T10.1. The blue team receives request for help at a distant location.
T11.1. Snake 2 is set to be the defector, and Snake1 is set to be the hero.
T12.1. Bandits turn to pursue Snake2.
T13.1. Snake1 fires two missiles at bandits.
T14.1. Bandits fire one missile each at Snake2.
T15.1. One of the bandits is destroyed.
T16.1. Snake2 is destroyed.
T17.1. Bandits fire one missile each at Snake1.
T18.1. Snake1 is destroyed.

Working narrative altered at T10 with character-actor separation (character swapping allowed):
T10.2. Snake 2 is set to be a defector, Snake1 is set to be a hero.
T11.2=T12.1. Bandits turn to pursue Snake2.
T12.2. Snake1 and Snake2 swap their characters. Snake1 proceeds to escape.
T13.2=T11. Snake2 fires two long-range missiles at bandits.
T14.2=T12. One of the two bandits is destroyed.
T15.2=T13. Snake2 illuminates the remaining bandit using the radar.
T16.2=T14. The bandit fires two missiles at Snake2.
T17.2. Snake 2 is destroyed, Snake1 is beyond the radar detection range.

5.2 Fruit collection

Two humanoid robots need to collect fruit (apples and oranges) scattered in equal numbers in a square room, placing them separately in two baskets (sorting fruits after collection is not allowed). Each robot can collect fruit with both hands, carry one basket and collect fruit with one hand, or carry two baskets. One of the robots is weak and unable to carry a heavy basket. A challenging situation occurs when this condition is discovered in the middle of task execution. The expected solution involves two characters: a basket carrier and a fruit collector.

Following the CR procedure, the actors will formulate several potentially useful characters: a basket carrier, a fruit collector (with a basket), and a fruit collector without a basket who will use both hands to collect fruit. The HNN will include the following states.

```
State S1: The room contains two robots, two baskets, and
          scattered apples and oranges.
State G1: All apples are in one basket.
State G2: All oranges are in one basket.
Goal = G1 and G2.
```

Then, among the generated narratives will be the following:

```
Narrative 1: S1 → (Arc1, Arc2) → Goal
      Character1: Apple collector. Performer: Robot1.
      Arc1: Goal=G1. Pick a basket, collect all apples. Reach G1.
      Character2: Orange collector. Performer: Robot2.
      Arc2: Goal=G2. Pick a basket, collect all oranges. Reach G2.
Narrative 2: State1 → (Arc3, Arc4) → Goal
      Character3: Basket carrier. Goal=G1 & G2. Performer: Robot1.
      Arc3: Pick baskets, follow the partner, carry baskets and
            ensure that each type of fruit is separated. Reach Goal.
      Character4: Fruit collector without a basket. Goal=G1 and G2.
      Performer: Robot2.
      Arc4: Collect apples and oranges, placing them separately in
            baskets carried by the partner. Reach Goal.
```

Narrative 1 is potentially more efficient than Narrative 2, and the actors can infer this from simulations. Indeed, suppose that each type of fruit is uniformly distributed in one of the two equal halves of the square environment, and the distributions do not overlap. Then Narrative 2 may take approximately twice the time required by Narrative 1. However, Narrative 1 cannot be continued when the weak robot has a heavy basket in its hands. When this condition is not known a priori and is discovered in the middle of task execution, CR in the weak actor will be triggered, resulting in its switching to Narrative 2. When the partner observes the weak robot approaching empty handed, it will put its behavior in the context of the two narratives, and will find the observed behavior consistent with Narrative 2, not with Narrative 1. Therefore, the robot-partner will infer that its partner decided to switch to Narrative 2, and will accept the new choice, even if the reason is not understood (the information about weakness may not be available to the partner). This example also illustrates how the global team GR is decomposed into local CR.

5.3 Bucket brigade

A firefighting team is distributed in a building affected by fire (Figure 3) where they have limited mobility. Because of the fire and building damage, actors are confined in local domains (rooms) and cannot easily cross obstacles separating them, but can pass buckets of water to each other through holes in the obstacles. Their task is to deliver water from the source to the fire sites, returning empty buckets back to the source of water. However, the actors cannot use verbal communications due to noise. Only simple gestures can be used to send two possible messages: (a) "do not send water here", (b) "send (more) water here". Messages can be sent to nearest neighbors only. The first message (a) is sent when a full bucket should be returned. It may indicate that all fires in that direction are extinguished, or that access to them is blocked. The second message (b) is initiated whenever a new fire is detected, and is transmitted along the chain toward a water source.

The team uses a fixed, small number of buckets. All buckets look distinct from each other and can be identified by their unique color. Actors have perfect episodic memory. In particular, they remember: which buckets they have passed and received; when, to/from whom, and in what condition; and when, from whom and in what number did the water requests come. Actors do not know the entire floor plan and the locations of fires, but remember how they got to their locations from the entrance (which is also the location of the water source; see Figure 3). Therefore, each actor knows where to send requests for water.

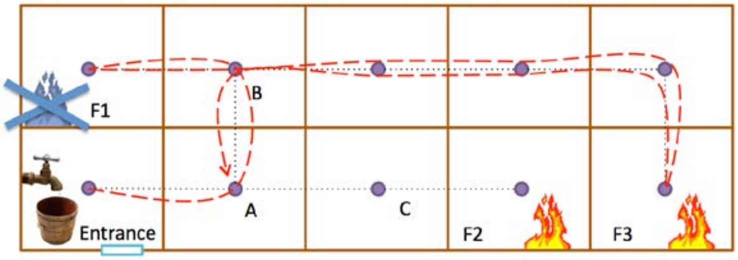

Figure 3. Floor plan of a building under fire with actors {A, B, C} and fires F1-F3. The water source is at the building's entrance. Dotted lines show possible routes of buckets and messages. The red dashed line shows the long path of the bucket given by A to B, after F1 was extinguished: B passes the bucket toward F1, but the bucket returns full, then it follows through B to F3 (cannot reach F2), back empty to B and to A.

Imagine the following scenario (Figure 3). At the beginning, requests for water arrive from F1, F2, and F3. Given that A received two requests from B and one from C, he will send twice as many buckets to B than to C. When F1 is extinguished, the cancellation message (a) will not reach A, because B will not pass the full bucket back to A: instead, he will pass it toward the remaining active fire F3. However, if A tracks individual buckets, then he will notice that it took an unusually long time for a bucket to return empty. Then A will interpret this event by simulating the bucket as a character, who returned from an extinguished fire and went to the other fire (A knows that there were two fires served by B and one served by C). This will explain the time increase. Therefore, A will start sending less water to B. In contrast, in a simple-minded approach without bucket tracking and without CR, A will not have a reason to change behavior.

5.4 Discreet pursuit

This scenario relates to the times when widespread surveillance was not in use. Suppose that a team of secret agents (actors), dressed like ordinary pedestrians, are monitoring a suspect in a city. Their goal is to watch the suspect continuously from a distance not exceeding the range of visibility (e.g., one block), in order to identify a secret location in the city, toward which the suspect is presumably heading. At the same time, their trajectories should be “believable” (not causing suspicion) from the suspect’s perspective. We also assume actors immediately see and recognize each other (or the suspect) within the range of visibility. They use wireless communication to broadcast messages to the team (e.g., a suspect’s coordinates). The suspect does not know the secret agents. His goal is to reach a secret location in the city without being followed, and he suspects that he might be followed. Therefore, he will change his direction randomly several times, and will not proceed to his destination if he observes that somebody followed his random turns without an obvious reason. In the initial state, one actor follows the suspect, while others are located nearby but outside of range. What strategy should the team use?

A CR solution can be formulated using three character types: a pedestrian, a shadow, and a double. A pedestrian trajectory should be directed toward some location in the city (this location may remain ambiguous, but should not change inconsistently or improbably) and should not follow the suspect for too long. The shadow (only one) stays within the range from the suspect. A double follows the suspect outside of the range (on a parallel street) based on broadcasts. The actor who follows the suspect plays the shadow and a pedestrian at the same time; other actors play doubles. The current shadow continues following the suspect until the two characters come into conflict with each other (e.g., when the suspect makes a random turn). At this moment, the actor transfers the shadow character to a double, and continues playing a pedestrian. It would be difficult to formulate a solution so concisely without using CR.

5.5 Black box recovery

The following scenario illustrates that CR can be useful not only in a team, but also in isolation. The actor is an unmanned underwater vehicle (UUV) that autonomously performs a task of search of a dead UUV. Upon locating it, the actor will go to the surface and communicate the location of the detected UUV to the operator (this operation takes some time). This task requires exploration of a large area of the seabed. The actor has already searched exhaustively one half of it, area A. The other half, area B, remains unexplored. At this point, the actor receives new information from its human operator: a plane has crashed somewhere in the same area. The black box (BB) will continue to emit signals for a month, and locating it is very important. Recovery of the UUV is equally important, but there is no short deadline associated with it. The two search tasks conflict with each other. A search for BB involves listening to a certain frequency of sound that can be heard at a greater distance compared to the distance of the UUV’s visibility. Therefore, an exhaustive BB-search can be performed faster, although it may not guarantee locating the UUV. The area A that is already searched for the UUV may contain BB. Assume that an exhaustive BB search of A and B will guarantee locating BB and will take up to 15+15 days. Similarly, assume that an exhaustive UUV search of B will guarantee locating the UUV and will also take 30 days. In responding to this situation, one approach could be to formulate possible goals and select one of them. In this case, the choice of the first goal would be to search A and B for BB, until it is found and its coordinates are communicated to the operator. However, the order in which A and B are searched remains ambiguous. Also, with this approach, if a UUV is spotted by chance during the BB search, it would be irrelevant to the current goal and will have no effect. An alternative approach for the actor is to play two characters in parallel: one is active (searching for a BB) while the other is dormant (searching for the UUV). The dormant character will bias decision making, attention, and behavior in cases when the active character has no preference. Thus, the actor will start its BB search from Area B, and will remember the location of a detected UUV to communicate it later (its immediate surfacing will disrupt the BB search).

6. Discussion

We presented a model of a Character Reasoner, which is combined with an actor serving as a team member. This means that one and the same cognitive system implemented in a robot needs to perform reasoning at multiple levels (Figure 2): (i) the author level, reasoning for the team and producing a whole-team fabula; (ii) the storyteller level, translating the fabula into a sjuzet: the arc of one character that will guide GR and planning from the local perspective; and (iii) the actor level, at which the system performs GR and planning guided by the sjuzet associated with one selected character. We illustrated this scheme in use and its benefits in five scenarios (none of them were implemented computationally: this is left for a future publication). Their comparison and specific points are summarized in Table 1, that also presents a comparison of CR with NTP. We see, in particular, that local decomposition of team-GR based on CR is possible in all cases.

Table 1. Comparison of the five scenarios. *Asterisks mark features that are unavailable in NTP.

	Sec.5.1	Sec. 5.2	Sec. 5.3	Sec. 5.4	Sec. 5.5
Feature or characteristic	Air combat	Fruit collection	Bucket brigade	Discreet pursuit	Black box recovery
Team-GR decomposed into local actor-GR		✔			
Team-GR decomposed into local CR*	✔	✔	✔	✔	✔
New characters *designed* in a new situation*	✔				✔
Characters associated with inanimate objects			✔		
Team or mission top goal selection	✔				✔
Narrative switching given a new situation	✔	✔	✔		✔
Character swapping during task execution*	✔	possible	✔	✔	
Character's top goal changes within an arc		✔	✔	✔	✔
Actor performs two characters in parallel*				✔	✔
Dormant character usage*					✔

We used the analyzed abstract examples to illustrate benefits of CR, point by point. Thus, in Section 5.2 (fruit collectors), the case of switching from Narrative 1 to Narrative 2 addresses the core question: is this model adaptive to changing environments or changing conditions? Our analysis of the Character Reasoner model gives an answer, explaining how and when the switching will occur, and why a global team-level GR may not be required. The analysis in Section 5.3 briefly conveyed that applying CR to objects (in this case, buckets) may benefit team performance. Section 5.4 further illustrates the capabilities of CR as opposed to NTP; we explain how one actor can play two characters simultaneously and benefit from this strategy. Section 5.5 extends this last point to a scenario involving only one actor in isolation, showing that the Character Reasoner model may be useful not only in teams or scenarios involving multi-agent interactions. Overall, CR is conceptually distinct from non-CR or non-narrative reasoning. It would be informative to compare them empirically using appropriate performance metrics; we expect that introducing characters in GR could be beneficial in some scenarios, as suggested by our example scenarios. In this case, the CR method will find applications in various domains. One possible example is the virtual football, the state of the art in which does not involve CR or NTP (Nadarajah & Sundaraj, 2013).

More generally, we would like our artificial actors to generate, understand and manage top-level goals in unpredicted situations. This level of cognitive autonomy is vital in scenarios involving multiple actors operating in dynamic unpredictable environments, especially when communications are limited. Scenarios usually involve goal-directed behavior that needs to be planned. In classical planning, an actor pursues a fixed goal. However, in real-world situations goals often need to be reprioritized, altered, or modified (Vattam et al., 2013). Autonomy in this case means that actors can perform GR both locally and consistently with their team, which is the focus of the CR framework we presented.

Young and his colleagues have described related work addressing GR in narrative generation (e.g., Riedl & Young 2010; Young et al., 2013). Their primary focus was to use automated planning to generate narratives and scripts for written or visual storytelling. Their developed NTP framework uses planning algorithms to produce desired narratives. As a result, the goals of individual characters are generated and managed during this process. Thus, planning can be used as a GR device to generate narratives.

7. Conclusions

- This paper introduced CR, a conceptual GR framework that uses narrative techniques to control the multi-actor behavior in multi-character scenarios. It extends the state-of-the-art on GR research, which previously did not focus on scenarios and solutions of this type (see Klenk, Molineaux, & Aha, 2013; Vattam et al., 2013, Samsonovich, 2014).
- Prior work related to this topic that has been studied in the literature on narrative intelligence has focused on planning (Riedl & Young, 2010; Young et al., 2013; Finlayson & Corman, 2013) and as such, has limitations. We have argued that the potential relative benefits of CR include overcoming limitations of the NTP approach identified earlier (Riedl & Young, 2010), and allowing a team of actors to use local GR instead of a global, team-level GR, thereby reducing the time and resources required for GR.
- We illustrated CR using example scenarios of cooperative problem solving. While we did not present empirical or formal analytical proofs in support of our claims, the presented informal analysis of selected examples strongly suggests that this topic is worthy of further study. Our

informal analyses of these scenarios suggests that reasoning in terms of characters differs from reasoning in terms of actors or narratives (when characters are not distinguished from actors), and that this may yield practical benefits in these and related scenarios. Future work suggestions include empirical computational studies using the formalism introduced here applied to examples similar to the analyzed scenarios.

Acknowledgements

Thanks to researchers collaborating with NRL's Adaptive Systems Section, especially Matt Molineaux, Ron Alford, Swaroop Vattam, and Mak Roberts, for their helpful discussions. The funding for this research was provided by OSD ASD (R&E). The views and opinions contained in this paper are those of the authors, and should not be interpreted as representing the official views or policies, either expressed or implied, of OSD or NRL.

References

Abell, P. (1987) *The Syntax of Social Life: The Theory and Method of Comparative Narratives*. Oxford University Press, Oxford.

Abell, P. (1993) Some aspects of narrative method. *Journal of Mathematical Sociology*, 18, 2-3, 93-134.

Abell, P. (2009) A Case for cases: Comparative narratives in sociological explanation. *Sociological Methods and Research*, 38(1):38-70.

Abell, P. (2011). Singular mechanisms and Bayesian narratives. In: Demeulenaere, Pierre, (ed.) *Analytical Sociology and Social Mechanisms*, pp. 121-135. Cambridge University Press, Cambridge, UK. ISBN 9780521154352.

Bal, M. (1998). *Narratology: An Introduction to the Theory of Narrative*. Toronto UP.

Finlayson, M. A., & Corman, S. R. (2013). The Military Interest in Narrative. *Sprache und Datenverarbeitung*, 37 (1-2).

Gray, W. D. (Ed.) (2007). *Integrated Models of Cognitive Systems. Series on Cognitive Models and Architectures.* Oxford, UK: Oxford University Press.

Haven, K. (2014). *Story Smart: Using the Science of Story to Persuade, Influence, Inspire, and Teach*. Santa Barbara, CA: ABC-CLIO, LLC. ISBN: 9781610698115.

Haven, K. (2007). *Story Proof: The Science Behind the Startling Power of Story.* Westport, Connecticut: Libraries Unlimited. ISBN 978-1-59158-546-6.

Huhn, P., Meister, J.C., Pier, J., and Schmid, W. (Eds.) (2013). *The Living Handbook of Narratology* (online). Hamburg: Hamburg University Press. URL: http://wikis.sub.uni-hamburg.de/lhn/

Kapadia, M., Falk, J., Zund, F., Marti, M., Sumner, B., and Gross, M. (February 27, 2015). Computer-assisted authoring of interactive narratives. *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (i3D 2015).*

Klenk, M., Molineaux, M., & Aha, D. W. (2013). Goal-driven autonomy for responding to unexpected events in strategy simulations. *Computational Intelligence*, 29(2), 187-206. doi: 10.1111/j.1467-8640.2012.00445.x

McCarthy, J. (1994). Situation calculus with concurrent events and narrative. Retrieved from http://jmc.stanford.edu/articles/narrative.html on 4.28.2015.

McCarthy, J. and Costello, T. (1998). Combining narratives. Retrieved from http://jmc.stanford.edu/articles/narrative2.html on 4.28.2015.

Nadarajah, S., & Sundaraj, K. (2013). A survey on team strategies in robot soccer: team strategies and role description. *Artificial Intelligence Review*, *40*(3), 271-304.

O'Neill, B. and Riedl, M.O. (2014). Dramatis: A Computational Model of Suspense. *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, Quebec City, Ontario.

Pentland, B. T., & Feldman, M. S. (2007). Narrative networks: Patterns of technology and organization. *Organization Science*, 18(5), 781-795. doi: 10.1287/orsc.1070.0283

Riedl, M.O., Stern, A., Dini, D., and Alderman, J. (2008). Dynamic experience management in virtual worlds for entertainment, education, and training. *International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning*, 3(1).

Riedl, M.O. & Young, R.M. (2010). Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research* 39: 217-268.

Riedl, M.O. & Young, R.M. 2014. The importance of narrative as an affective instructional strategy. In R. Sottilare, A. Graesser, X. Hu, and B. Goldberg (Eds.). *Design Recommendations for Adaptive Intelligent Tutoring Systems: Adaptive Instructional Strategies*, volume 2. Army Science Laboratory.

Roberts, M., Vattam, S., Aha, D.W., Wilson, M., Apker, T., & Auslander, B. (2014). Iterative goal refinement for robotics. In: A. Finzi & A. Orlandini (Eds.) Planning and Robotics: Papers from the ICAPS Workshop. Portsmouth, NH: AAAI Press (in press).

Samsonovich, A. V. (2014). Goal reasoning as a general form of metacognition in BICA. Biologically Inspired Cognitive Architectures, 9: 105-122. DOI: 10.1016/j.bica.2014.07.003.

Schmid, W. (2010). *Narratology: An Introduction*. Walter de Gruyter GmbH & Co. KG, Berlin/New York. ISBN 978-3-11-022631-7.

Simon-Shoshan, M. (2013). *Stories of the Law: Narrative Discourse and the Construction of Authority in the Mishnah.* Oxford: Oxford University Press.

Vattam, S., Klenk, M., Molineaux, M., & Aha, D.W. (2013). Breadth of approaches to goal reasoning: A research survey. In D.W. Aha, M.T. Cox, & H. Muñoz-Avila (Eds.) Goal Reasoning: Papers from the ACS Workshop (Technical Report CS-TR-5029). College Park, MD: University of Maryland, Department of Computer Science.

Ware, S.G., and Young, R.M. (2014). Glaive: A state-space narrative planner supporting intentionality and conflict, in *Proceedings of the 10th Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE2014)*, pp. 80-86. Raleigh, NC.

Young, R.M., Ware, S.G., Cassell, B.A., and Robertson, J. (2013). Plans and planning in narrative generation: A review of plan-based approaches to the generation of story, discourse, and interactivity in narratives. *SDV: Sprache und Datenverarbeitung: Intenational Journal of Language Processing*, 37 (1-2): 41-64.

Zimmerman, B.J. (2002). Becoming a self-regulated learner: An overview. *Theory into Practice*, *41*(2): 64-70.