Motivation Discrepancies for Rebel Agents: Towards a Framework for Case-based Goal-Driven Autonomy for Character Believability

Alexandra Coman[1], Héctor Muñoz-Avila[2]

[1] Department of Electrical and Computer Engineering and Computer Science, Ohio Northern University, Ada, OH 45810
[2] Department of Computer Science and Engineering, 19 Memorial Drive West, Lehigh University, Bethlehem, PA 18015
a-coman@onu.edu, hem4@lehigh.edu

Abstract. Character believability is one of the key requirements of a successful narrative. In order to be believable, characters must appear to act in accordance with personal motivation, which is shaped through events occurring throughout the narrative. Goal-Driven Autonomy (GDA), is a model of goal reasoning in which agents identify, reason about, and, if necessary, modify goals that they will pursue. Our proposal for using GDA in character-centric interactive narratives is based on the new concept of "rebel agents" and on a set of new discrepancies, which we call "motivation discrepancies". A rebel agent is an agent that may "refuse" a goal, plan, or plan component that it assesses to be in a conflict with its own motivation. Previously, in systems performing planning and execution, including GDA agents, a plan assigned to an agent was executed unless exterior conditions made it unfeasible or no longer useful to do so. In the case of rebel agents, plans may fail to be executed and new goals may be chosen due to the agent's own "refusal", based on the agent's internal motivation. Motivation discrepancies cover a wider range of situations than the traditional GDA discrepancies, and are suited specifically to character-believability purposes.

Keywords: goal-driven autonomy, character believability, interactive storytelling

1 Introduction

Character believability (Bates, 1994) is considered to be one of the key requirements of a successful narrative, be it interactive (e.g., a computer game) or traditional (e.g., narratives in novels). Believable characters act in accordance with personal memories and motivations, which are shaped through events occurring throughout the narrative. Motivation and memories should evolve as the story progresses, so as to create plausible and engaging character growth (Schaefer & Diamond, 1998).

Goal-Driven Autonomy (GDA) is a model of goal reasoning in which agents identify, reason about, and, if necessary, modify goals that they will pursue (Muñoz-Avila et al., 2010a; Aha et al., 2011). Goal modifications are undertaken as a result of identifying discrepancies between expected and actual states of the environment during plan execution (Cox, 2007). The potential for GDA has been demonstrated in a varie-

ty of domains, including real-time strategy games (Weber et al., 2010; Jaidee et al. 2013) and navy simulations (Molineaux et al., 2010).

We propose an extended form of GDA that enables creating believable characters, and can be used, for example, in interactive narrative worlds structured as multi-agent systems coordinated by a "drama manager". Multi-agent systems are a particularly good fit for GDA as each character can be controlled by its own agent, and the overall coordination or storyline can be controlled by another agent (Jaidee et al., 2013).

We propose to adapt the GDA framework to character-centric narrative generation by introducing new types of GDA discrepancies, which we call motivation discrepancies. Thus, the notion of discrepancy no longer refers strictly to mismatches between expected states and actual states, but includes more subtle incongruities, such as those between a character's changed motivation and the character's previously-assigned goal/course of action.

2 Goal-Driven Autonomy

GDA agents integrate planning and execution, and introspectively analyze their own course of action to determine if and when it is necessary to pursue new goals. The basic GDA cycle works as follows:

(a) The agent generates a plan π for a given goal g. When an episode starts, the first goal can be a default one or one assigned by the user.
(b) The agent executes the actions in the plan π in the environment.
(c) The agent monitors the plan's execution: after executing each action a in π, the agent checks if the resulting state s matches the expectations x of action a.
(d) If there is a mismatch between the expectations x and the actual state s, a discrepancy d has been found.
(e) The agent generates an explanation e for the discrepancy d.
(f) Based on the discrepancy d, the actual state s, and the expectation x, the agent either keeps the same goal or generates a new goal g'.[1] As an example of the former, the game-playing GDA agent reported in Jaidee et al. (2013) keeps the same goal when it is winning the game. The basic idea is that even though action(s) in the plan π are not fulfilling the desired expectations, pursuing the current goal g is enabling the agent to win. Hence, the agent sees no reason to change its goal.

The designer of the agent must consider a number of issues, including: (1) when to check if the effects of an action are fulfilled (for example, the agent might check the state of the environment after some pre-defined period of time), (2) whether to pursue multiple goals at the same time, and (3) even if it is pursuing one goal at a time, establishing a ranking to determine which goal should be pursued next.

[1] We are assuming an agent that pursues one goal at a time. However, GDA agents can pursue multiple goals at a time. For such agents, the new goal g' is simply added to the list of goals to achieve.

3 Case-based Goal-Driven Autonomy Agents

CBR is a natural choice for building GDA agents because it can significantly alleviate the knowledge-engineering requirement to build such agents (e.g., Muñoz-Avila et al., 2010b). Specifically, there are three pieces of GDA knowledge that can naturally be represented as cases.

First, the planning knowledge can be represented as a library of (g, π) cases, which associate with each goal g a plan π that achieves it. Furthermore, similarity criteria enable the reuse of cases when the goals of the current situation are similar to cases' goals (as opposed to requiring the goals to be identical). Second, the expectations knowledge base can be represented as a case base of triples (s,a,x), where s is the current state, a is an action, and x is the expected state to be obtained after executing a. Once again, similarity criteria enable GDA agents to relax the notion of expectations: observed states may be acceptable if they are similar to expected ones, even though not identical. Third, the next goals to pursue can be represented as a case base of triples (d,g,g') where d is the discrepancy, g is the current goal and g' is the next goal to pursue. Again, similarity criteria on the states and goals can make the cases applicable in a wide range of situations.

4 Goal-Driven Autonomy for Believable Agents

We believe that GDA is a natural fit for character believability for the following two main reasons:

(1) It is based on agent autonomy and decision power: a character that has the capability of changing its goals based on reasons it identifies by itself, as the story progresses, manifests intentionality, a key component of character believability (Riedl and Young, 2010).
(2) The notion of discrepancy, which is central to GDA, is also the driving force of conflict in many narratives. Discrepancies such as those between a hero/heroine's will and his/her destiny or the will of the gods; love and duty, loyalty to family and allegiance to liege lord, etc. create narrative tension.

In previous systems performing planning and execution, including GDA agents, a plan assigned to an agent was executed unless exterior conditions made it unfeasible or no longer useful to do so. What is different in our proposed work is that plans may fail to be executed and new goals may be chosen due to the agent's own "refusal" based on its internal motivation, where this motivation is decoupled from the plan itself.

Traditional GDA discrepancies are created by external factors, i.e. they are discrepancies between the actual state and the expected state according to the action model. Our proposed discrepancies are between external and internal factors, e.g. between an observed state and the internal motivation of the agent; crucially, this motivation is the agent's own, and not something dictated by an action model.

Our proposal for using GDA in character-centric interactive narratives is based on the new concept of the **rebel agent** and on a set of new discrepancies, which we call **motivation discrepancies**:

(a) The Rebel Agent is an agent that may "refuse" a goal, plan, or plan component that it assesses to be in a conflict with its own motivation. This type of agent is not restricted to goal-driven autonomy, but could be part of other types of planning/reasoning models.
(b) Motivation discrepancies are discrepancies that cover a wider range of situations than the traditional GDA discrepancies, and are suited specifically to character-believability purposes.

5 Related Work

Goal-driven autonomy was in part motivated by work on self-aware agents (Cox, 2007). The first publications on this topic present the basic cycle described in the previous section (Muñoz-Avila et al., 2010a; Molineaux et al., 2010). This initial work frames GDA as tightly coupling deliberative planning and execution, and distinguishes GDA from pure deliberative planning approaches (Ghallab et al., 2004) and pure reactive planning approaches, where the plan needs to be adjusted but the goals remain the same (Firby, 1987).

To ensure consistency in the behavior of non-player characters (NPCs), many commercial digital games assign a faction to NPCs in the game a priori. As the player's avatar progresses through the game, it may join and leave factions. The game pre-establishes rules about relations between factions. For example, factions might be enemies, indifferent, or friendly to one another. As a result, the player will be treated consistently by all NPCs from the same faction, thereby ensuring some level of believability in the characters (Sweetser & Wyeth, 2005). Dinerstein et al. (2007) use reinforcement learning techniques to create agents that mimic humans, with the goal that the agents behave more realistically. Other authors see the role of NPCs as actors and include mechanisms to ensure that the behavior is consistent with the overarching storyline (Magerko et al., 2004). AI planning techniques have also been used to ensure that the NPC's intent is consistent with the overarching storyline (Riedl and Young, 2010). In this sense, Riedl and Young follow what they call deliberative narratives, as the AI planner is generating the entire storyline, including the actions of all characters. This, as opposed to the more traditional emergent narratives were the system is a sandbox and the players' actions determine the storyline (Aylett, 1999).

The closest work to ours we have seen is the AI for a commercial game series called Close Combat. This is a real-time strategy game in which the player controls a small number of groups or "squads" of characters. Each of these squads has associated a number of numerical factors such as stress (i.e., a degree of combat squad the unit experiences; if the squad is not in combat, then its stress level is at its lowest value). These factors are combined using fuzzy logic operators (Sweetser & Wiles, 2002). However, unlike in our research, these are not autonomous agents. They follow orders until their level of stress exceeds a certain threshold.

6 The Rebel Agent

We define a "rebel agent" as one being capable of "refusing" to perform a goal or plan assigned by a human user or by another AI agent. The rejection is the result of conflict between the goal or plan and the agent's own internal motivation, modeled for the purpose of creating character believability. By "rebel agents" we, of course, do not mean agents the behavior of which has been hardcoded or created by a central story planner so as to appear rebellious in the context of a specific given storyline. Instead, the refusal is based on the agent's own internal motivation. We hypothesize that this will result in an increase in the believability of the character's behavior.

The rebel agents' course of action following a goal or plan refusal may vary. Two possible reactions are (1) choosing a new goal, as in GDA, and (2) replanning. It follows that the notion of rebel agent is not restricted GDA.

What distinguishes rebel agents from other agents (which we could call "compliant agents") that can conduct goal reasoning and replanning is the fact that, in the case of rebel agents, goal/plan changes are undertaken specifically as a result of the agent "taking issue" with the plan/goal itself, which it "perceives" as coming into conflict with its own motivation.

We now describe a possible context in which we envision such an agent model being used: an interactive narrative world in which a central "drama manager" (itself possibly playing a narrative role, such as that of "destiny" or "commander") assigns plans to agents, possibly "in bulk" (e.g. all "soldier" characters receive the same plan). This would roughly fall under the category of "emergent system" coordinated by a "drama manager", using the terminology of Riedl and Young (2010). Another important characteristic of the set-up is that it should make sense, for believability purposes, for centrally-assigned plans to be "rejected" by an individual agent under particular circumstances that the agent may encounter (e.g., the command is to attack the nearest enemy, but, for one of the agents, the nearest enemy is in a civilian location, and causing harm to civilians is against the agent's own beliefs). It should be noted that, in an actual implementation, human players could fulfill one of these roles in the system. For example, the player could interfere with the story by modifying the motivation of one or more agents, thus aligning with the rebel agents and contributing to the tension between the drama manager and the agents. Alternatively, the player could take on the role of the drama manager, and be responsible for giving commands which may or may not be obeyed.

As an example, consider an environment in which we have fighting units of various types (e.g. infantry and cavalry): all infantry characters receive the command to attack a given location, i.e. they are assigned an attack plan. Most soldiers abide by the plan. One soldier, however, refuses to do so outright, because his motivation to remain by his family is stronger. Another soldier carries out the plan up to a given point, but encountering his father on the enemy side eventually leads him to re-evaluate his goals.

Here is how responsibilities in terms of planning and knowledge-engineering effort might be shared between rebel agents and the drama manager:

- The drama manager is responsible for assigning goals/plans to agents. Plans may, for example, be assigned based on unit type. They are not based on the characters' evolving motivation, to which the drama manager is "blind".

- The agent is responsible for maintaining and updating its own internal motivation.

The generation of a new goal may be conducted at agent level (i.e. following the GDA cycle), or it may be a collaboration between the agent and the drama manager: the agent might request that the drama manager assign it a new goal based on the discrepancy, or the agent might identify the appropriate new goal based on the discrepancy, and present it to the drama manager, so that a plan might be generated for it.

Whether the agent will "rebel" or not may be decided based on a measure of compliance. Compliance can be recalculated periodically or before specific events, such as executing an action. If the compliance threshold is met (e.g. the contrast between the action that the agent is about to execute and the agent's motivation is not evaluated as strong enough to justify rebellion) the assigned plan will continue to be carried out as it is. This is similar to the use of thresholds for inverse trust by Floyd, Drinkwater, and Aha (2014).

GDA explanations of discrepancies can be embedded in the storytelling through text or other visual cues in order to highlight the intentionality of the characters' actions, a crucial component of character believability (Riedl and Young, 2010).

Case-based reasoning techniques can be used, as previously in GDA, to identify new goals in mismatch situations. However, given the potentially higher complexity of motivation-based discrepancies and believable-character models in general, case-based reasoning, given its roots in human reasoning, could have a more extensive use throughout the process, e.g. in establishing motivation and identifying corresponding "personal" mismatches for the various characters. Its use could, therefore, be central in the motivation model upon which discrepancies are to be based, which we discuss next.

The Motivation Model. We have so far considered the character's motivation to be a black box, having only mentioned that it evolves, and that it does so in a decentralized manner, with changes being managed by the agent itself, not by the drama manager. The agent keeps track of its own motivation progress. For example, when meeting a character C that surpasses a certain level of social compatibility with agent A, agent A might update its internal status from not being in love to being in love with character C.

In order to model motivation, we could use one of or a combination of several of the pre-existing character modeling techniques, such as memory models, emotion models, character personality traits, and models based on social relationships (Brom, Peskova, & Lukavsky, 2007; Dias et al., 2007; Gomes & Jhala 2013; Gomes, Martinho & Paiva, 2011; Kope, Rose & Katchabaw, 2013; McCoy et al., 2010; Strong et al., 2007, and many others). The techniques used for the implementation of the motivation model might include ontologies (already successfully used in combination with GDA by Dannenhauer and Muñoz-Avila, 2013) and constraint satisfaction. The motivation model could also be handled using CBR techniques.

7 Motivation Discrepancies

In GDA research, discrepancies are defined as mismatches between the expected state and the actual observed state. In practice, both of these states have so far described circumstances external to the agent, that is, conditions that are observable in the state of the world (e.g., in the gaming environment); there was, so far, no notion of internal agent motivation in the GDA model.

For the purposes of character believability, we propose the new kinds of discrepancies listed below, which we call motivation discrepancies. For each of type of discrepancy, we also briefly explain why it is not reducible to the discrepancy between expected state and observed state.

(a) **Discrepancy between motivation and the next plan action:** The agent "refuses" to execute the next action in the plan because the action itself or the combination of the action and its parameters contravene(s) the agent's motivation. For example, agent Juliet is assigned the plan "<attend(Ball)>, <marry(Paris)>". Her motivation model, based on social relationships, contains the fact inLoveWith(Romeo). She executes the first action in the plan, but refuses the next one, and requests a new goal, compatible with her motivation (e.g. elopement with Romeo). This type of discrepancy is not reducible to the traditional discrepancy because, sometimes, it makes sense for a character to reject an action as soon it is encountered in the plan, before executing it and "seeing" what its consequences are. Because Juliet is in love with Romeo, she will immediately reject the action "Marry Paris", without waiting to observe its consequences as expressed in the resulting state. Note that, when the plan was initially assigned, Juliet may not have met Romeo yet, so the plan may have been acceptable to her at that time.

(b) **Discrepancy between motivation and observed state**: The agent adjusts its goal because it finds a discrepancy between an observed state and its motivation model. For example, agent Clytemnestra, while pursuing a goal assigned to her by King Agamemnon, executes a plan containing the action <bringTo(Iphigenia, Aulis)>. After executing the action, she finds, in the current state, that her daughter Iphigenia has been sacrificed to appease the goddess Artemis. This observed state is so starkly in conflict with her internal motivation (which is dominated by maternal love), that she modifies her goal to that of revenge against Agamemnon. This type of discrepancy is not reducible to the traditional discrepancy because motivation is not part of the expected state.

(c) **Discrepancy between motivation and current goal:** A goal might be rejected and exchanged straight away, without even initiating the execution of its associated plan, if it is deemed incompatible with the agent's current motivation. As her motivation model includes the character trait "ambition", Agent Macbeth might immediately identify a goal involving subservience as unacceptable. This would occur before plan execution, so before observing any states resulting from plan execution, which is why this type of discrepancy is also not reducible to the traditional one.

Of course, regular GDA can be used for the purpose of creating character believability without introducing new types of discrepancies: for example, by making the states include internal agent motivation characteristics and by designing the structure of the case library accordingly. However, we assume a model in which the character's motivation and the story action, though interrelated, proceed in a decoupled manner, so as to ease the burden on the drama manager and to allow asynchronous character development. Also, we believe that extending the notion of discrepancies to include motivation discrepancies is a more natural approach to GDA for this domain, much like creating a type of planner specific to story-telling was found by Riedl and Young (2010) to be a better solution for the narrative-generation problem than simply using traditional planners with narrative-specific domain descriptions.

8 Conclusions and Future Work

In this paper, we define and explore general characteristics of "rebel agents". The concept of rebel agent involves an apparent contradiction: we want the agent to generally follow plans and goals assigned by the player or the drama manager, but we also want it to have the potential of being rebellious and refusing to follow the instructions under certain circumstances. Furthermore, we want the agent's behavior to be believable, so rejection must be consistent with the characters' previous and future behavior, experience, and development.

Our solution is to use GDA and to assign the agent a motivation model.

We introduce motivation discrepancies, which are discrepancies between what the agent is instructed to do and the agent's own motivation. When deemed appropriate, the agent will autonomously rebel against the order and follow new goals that are consistent with its own motivation.

For future work, we want to implement rebel agents and test them. We have several ideas for possible benchmarks: e.g. a GDA agent that does not have a motivation model (i.e., an ablation of the agent's rebellion) and a GDA agent that randomly (with a low probability) decides not to follow instructions. Potential performance metrics could include scores measuring to what extent what the agent achieves is consistent with the agent's motivation. For example, if the motivation of the agent is to minimize civilian casualties, while the objective is to maximize the score, which for example, increases as a function of the number of enemy units destroyed, then we could comparatively assess the agent's performance as a function of the score and the number of civilians killed. Ideally, the rebel agent will have the lowest number of civilians killed while still having a high score compared to the baselines.

The next stage of the experimental evaluation will involve asking users to assess the characters' believability.

Acknowledgements. We thank the reviewers for their comments and suggestions, some of which have been integrated into the final version of the paper, others of which will be of help to us in our future work. This work is supported in part by NSF grant 1217888.

References

Aha, D.W., Molineaux, M., & Klenk, M. (2011). Goal-driven autonomy. 2011 NRL Review, 164-166

Aylett, R. (1999). Narrative in virtual environments — towards emergent narrative. In

Bates, J. (1994). The role of emotion in believable agents. Communications of the ACM, 37(7), 122-125.

Brom, C., Peskova, K., Lukavsky, J., (2007). What Does Your Actor Remember? Towards Characters with a Full Episodic Memory. *ICVS 07: Proceedings of the 4th International Conference on Virtual Storytelling*, pp. 89-101.

Cox, M. T. (2007). Perpetual self-aware cognitive agents. AI Magazine 28(1), 32-45.

Dannenhauer, D. and Muñoz-Avila, H. (2013) LUIGi: A Goal-Driven Autonomy Agent Reasoning with Ontologies. Advances in Cognitive Systems Conference (ACS-13).

Dias, J., Ho, W., Vogt, T., Beckerman, N., Paiva, A., & Andre, E. (2007). I Know What I Did Last Summer: Autobiographic Memory in Synthetic Characters. *Affective Computing and Intelligent Interaction*, 606-617.

Dinerstein, J., Egbert, P. K., & Ventura, D. (2007, January). Learning Policies for Embodied Virtual Agents through Demonstration. *In Proceedings of the International Joint Conference on AI (IJCAI-07)*. (pp. 1257-1262).

Floyd, M.W., Drinkwater, M., & Aha, D.W. (2014). How much do you trust me? Learning a case-based model of inverse trust. To appear in *Proceedings of the Twenty-Second International Conference on Case-Based Reasoning*. Cork, Ireland: Springer.

Firby, R. James. (1987) An investigation into reactive planning in complex domains. In AAAI, vol. 87, pp. 202-206.

Ghallab, M., Nau, D., & Traverso, P. (2004). Automated planning: theory & practice. Elsevier.

Gomes, P., & Jhala, A. (2013). AI Authoring for Virtual Characters in Conflict. *Proceedings of the Ninth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pp. 135-141.

Gomes, P. F., Martinho, C., & Paiva, A. (2011). I've Been Here Before! Location and Appraisal in Memory Retrieval. *Int. Conf. on Autonomous Agents and Multiagent Systems*.

Floyd, M. W., Drinkwater, M., & Aha, D. W. (2014). Case-Based Behavior Adaptation Using an Inverse Trust Metric.

Jaidee, U., Muñoz-Avila, H., & Aha, D.W. (2013) Case-Based Goal-Driven Coordination of Multiple Learning Agents. the twenty-first International Conference on Case-Based Reasoning (ICCBR). Saratoga, NY.

Kope, A., Rose, C., & Katchabaw, M. (2013). Modeling Autobiographical Memory for Believable Agents. *Proceedings of the Ninth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pp. 23-29.

Magerko, B., Laird, J., Assanie, M., Kerfoot, A., & Stokes, D. (2004). AI characters and directors for interactive computer games. Ann Arbor, 1001(48), 109-2110.

Mateas, M., & Sengers, P. (Eds.), Narrative Intelligence: Papers from the AAAI Fall

Symposium (Technical Report FS-99-01), pp. 83–86. AAAI Press, Menlo Park.

McCoy, J.; Treanor, M.; Samuel, B.; Tearse, B.; Mateas, M.; & Wardrip-Fruin, N. (2010). The Prom: An Example of Socially-Oriented Gameplay. In Proceedings of the Sixth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE 2010), 221-222. AAAI Press.

Molineaux, M., Klenk, M., & Aha, D.W. (2010). Goal-driven autonomy in a Navy strategy simulation. In Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence. Atlanta, GA: AAAI Press.

Muñoz-Avila, H., Aha, D.W., Jaidee, U., Klenk, M., & Molineaux, M. (2010a). Applying goal directed autonomy to a team shooter game. Proceedings of the Twenty-Third Florida Artificial Intelligence Research Society Conference. Daytona Beach, FL: AAAI Press.

Muñoz-Avila, H., Aha, D.W., Jaidee, U., Carter, Elizabeth. (2010b) Goal-Driven Autonomy with Case-Based Reasoning. Proceedings of the 18th International Conference on Case Based Reasoning (ICCBR 2010). Springer.

Riedl, Mark O. & Young, R. Michael. Narrative Planning: Balancing Plot and Character. Journal of Artificial Intelligence Research, vol. 39, 2010.

Schaefer, C., & Diamond, R. (1998). The creative writing guide: Poetry, literary nonfiction, fiction, and drama. Longman.

Strong, C.R.; Mehta, M.; Mishra, K.; Jones, A.; and Ram, A. 2007. Emotionally Driven Natural Language Generation for Personality Rich Characters in Interactive Games. In Proceedings of the Third Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE 2007), 98-100. AAAI Press.

Sweetser, P., & Wiles, J. (2002). Current AI in Games: A review. Australian Journal of Intelligent Information Processing Systems, 8(1), 24-42.

Sweetser, P., & Wyeth, P. (2005). GameFlow: a model for evaluating player enjoyment in games. Computers in Entertainment (CIE), 3(3), 3-3.

Weber, B., Mateas, M., & Jhala, A. (2010). Applying goal-driven autonomy to StarCraft. In Proceedings of the Sixth Conference on Artificial Intelligence and Interactive Digital Entertainment. Stanford, CA: AAAI Press.