



Adaptation Vectors for Instructional Agents

Benjamin Bell^{1(✉)} and Robert Sottolare²

¹ Eduworks Corporation, Corvallis, OR 97333, USA
benjamin.bell@eduworks.com

² Soar Technology, Inc., Orlando, FL 32817, USA
bob.sottolare@soartech.com

Abstract. This paper examines the application of intelligent agents to guide and adapt instruction in a class of learning technologies known as adaptive instructional systems (AISs). AISs are *artificially-intelligent, computer-based systems that guide learning experiences by tailoring instruction and recommendations based on the goals, needs, and preferences of each individual learner or team in the context of domain learning objectives*. Intelligent agents are autonomous entities which observe the conditions in their environments through percepts (e.g., sensors) and then act upon the environment using actuators. Their activity is intelligently directed to strategies (plans for action) and tactics (actions executed by the AIS) which enhance the progress of the learner toward the achievement of assigned goals and objectives. To optimize agent and learner performance we examine a notional set of principles that is needed to guide adaptive instructional designers and system developers toward *adaptation vectors* that support both instructionally meaningful (effective) and doctrinally correct (relevant and tactically plausible) actions by the AIS. In conclusion, we explore the literature to understand the design of adaptation vectors and provide recommendations for future AIS research and standards development.

Keywords: Adaptation · Vectors · Adaptive instructional system (AIS) · Instructional strategies · Instructional tactics

1 Introduction

The focus of this paper is on the application of intelligent agents and their role in various tailoring methods for adaptive instructional systems (AISs). We have termed these tailoring methods *adaptation vectors*. We begin by defining what AISs and intelligent agents along with their major components and functions. We follow this up with a problem definition.

1.1 Defining Adaptive Instructional Systems

During the last year, an IEEE working group chartered through Project 2247 has taken on the task of developing standards and best practices for AISs. This IEEE working group will be debating what is and is not an AIS, and what standards and best practices will evolve from the marketplace. To date, the group has identified three potential areas

for standardization: (1) a conceptual model for AISs, (2) interoperability standards for AISs, and (3) evaluation best practices for AISs. This paper explores design principles and methods to enhance interoperability and reuse for AISs) that are defined as: *artificially-intelligent, computer-based systems that guide learning experiences by tailoring instruction and recommendations based on the goals, needs, and preferences of each individual learner or team in the context of domain learning objectives* [1]. It is important to distinguish between adaptive and adaptable system attributes. Both adaptive and adaptable systems provide system flexibility, but adaptive systems are able to observe the environment, identify changing conditions, and then take action without human intervention while adaptable systems provide control over change/flexibility to the user [2, 3]. Adaptive and adaptable attributes are not mutually exclusive and both adaptive and adaptable features can exist together in the same system. For example, the Generalized Intelligent Framework for Tutoring (GIFT) has an authoring system that features adaptability by allowing users to configure course objects and also features adaptivity by automatically building a tutor based on specifications developed by the author [4, 5].

Adaptive instructional systems (AISs) come in many forms and this makes standardization challenging. The most common form of AIS is the intelligent tutoring system (ITS) which is a computer-based system that automatically provides real-time, tailored, and relevant feedback and instruction to learners [6, 7]. Other forms of AISs include intelligent mentors (recommender systems) which promote the social relationship between learners and intelligent agents [8] and web-based intelligent media used for instruction. Next, we define intelligent agents and discuss their various attributes and forms.

1.2 Defining Intelligent Agents

Intelligent agents are autonomous entities which observe the conditions in their environments through percepts (e.g., sensors) and then act upon the environment using actuators [9]. Agents may take different forms:

- simple reflex agents [9] – act only in response to the current percept, ignoring the rest of the percept history; uses condition-action rules to determine action selection
- model-based reflex agents [9] – acts based on a model that represents percept history and interprets the current state of the world, how the world has evolved (trends), and the impact of decisions in terms of rewards; uses condition-action rules to determine action selection
- goal-based agents [9] – expand on the capabilities of model-based agents by using goal information that describes desirable states; uses progress toward goals to determine action selection
- utility-based agents [9] – define a measure of desirability of a particular state called a *utility function* which maps each state to a measure of the utility; improves upon goal-based agents which can only distinguish between goal states (desirable) and non-goal states (not desirable); acts to be in states of the highest utility
- learning agents [9] – allows the agents to initially operate in unknown environments and to become more competent through experiences; implements separate processes

for *learning* and *performance* which are responsible for making improvements and selecting external actions respectively

- multi-agent architectures [10] – composed of hierarchies of agents with different functions; the GIFT architecture [4, 5] is a multi-agent architecture with agents assigned to manage dialogue between learners and virtual humans, and decisions for recommendations, strategy selection, and tactic selection

A set of desirable characteristics for intelligent agents (derived from Russell and Norvig [9]) might include:

- Autonomous – agents are able to act without human direction or intervention
- Adaptive and Reactive – agents are responsive to changes in their environment and are active in enforcing policies (rules)
- Proactive – agents take initiative to achieve long-term goals, recognize opportunities to progress toward goals, and learn and adapt to make more effective choices in the future
- Sociable and Cooperative – in multi-agent architectures, agents share information and act together to recognize and achieve long-term goals

For AIS architectures, intelligent agents not only act on the environment, but also observe and act upon the learner(s) as shown in Fig. 1.

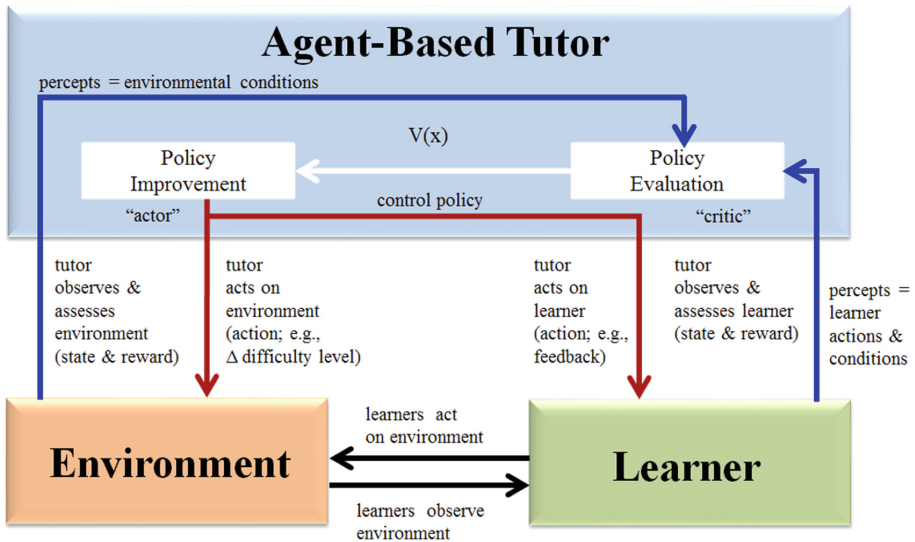


Fig. 1. Intelligent agents interacting with learners and the environment [11]

In AISs, the activities of effective agents are intelligently directed to develop recommendations, strategies (plans for action) and tactics (actions executed by the AIS) which enhance the progress of the learner toward the achievement of assigned goals and objectives. Policies (often rule-based, decision trees or Markovian processes) are

developed to aid AIS decision-making to select optimal next actions. The goal of each action is to optimize learning outcomes and policies drive recommendation, strategies, and tactics that have the highest rewards with respect to learning, performance, retention, and transfer of skills. Next, with a greater understanding of AISs and intelligent agents, we begin to shape our problem space.

1.3 Defining Our Problem Space

Much of research in AISs has been focused on blending agent-enabled instruction with adaptive tutoring. At the nexus of these converging areas is the construct of an adaptive instructional agent, a label that we use fully aware of the ambiguity between the adaptivity of the agent and tailoring of the instruction facilitated by an adaptive agent. In real-time, AISs use adaptive agents to autonomously make instructional decisions and as part of this process act on both learners and the environment.

To be effective in selecting actions, the adaptive agent must be cognizant of its environment and have a model of how its decisions will impact learning. Tailoring based on individual needs, goals, and preferences are dependent on a comprehensive learner model that includes learner attributes which are known to moderate or influence learning. So why are we staking out such a subtle distinction? The reason for this distinction is more than a semantic exercise and really has to do with the authoring process for adaptive instruction and specifically agent impacts on the environment (e.g., simulation, serious games, problem banks).

Research Question: How might we use intelligent agents to author more relevant and appropriate content to offer better, more diverse options for adaptive instruction – especially during remediation where new experiences are used to assist learners in achieving expected competencies after initial instruction results in learning gaps?

As we noted previously, intelligent agents come in different forms. Assuming we need agents that learn and are flexible, cooperative, proactive, and autonomous, we can envision agent policies that not only target a diversity of solutions, but are also provide doctrinally correct solutions. It is also important that agents are situationally-aware and have fitness criteria to weigh their decisions when choosing between multiple *good* actions that assist in progress toward multiple learning goals. All this complexity indicates that we should consider a multi-agent architecture, but where should we begin?

Agents that exhibit adaptive behaviors might be valuable in training simulations and educational courses by creating a variety of diverse challenges for learners and resulting in a variety of performance outcomes. While generating a variety of behaviors for constructive entities, sundry scenarios or an assortment of problem sets is an important element of adaptive instruction, by itself it is not enough. The agents operating as part of AISs or in partnership with AISs need to act using tactically-plausible adaptive behaviors, produce doctrinally relevant scenarios or produce rich, germane problems sets. Diversity is not enough. When we want intelligent agents to contribute to learning during instruction, the adaptive behaviors they exhibit must meet different standards and their policies must weigh outcomes in terms of their importance or fitness with respect to learning goals.

A policy (set of principles) is needed that can guide training developers in creating adaptive instructional agents for simulation-based training whose adaptations are both instructionally meaningful (effective) and doctrinally correct (relevant and tactically plausible). We term these *adaptation vectors*. In this paper, we explore what the adaptive tutoring literature has to offer the designers of agent-based instruction, discuss recent examples of adaptive, agent-based tutoring systems, and synthesize preliminary adaptation vectors from these case studies.

2 Adaptive Instruction vs. Adaptive Agents

Earlier we defined and listed essential attributes of intelligent agents and suggested a difference between an adaptive instruction and adaptive agents. Here we begin to compare and contrast these terms, and ultimately define an *adaptive instructional agent* and its essential policies. We define adaptive instruction as *guided learning experiences in which an artificially-intelligent, computer-based system tailors instruction and recommendations based on the goals, needs, and preferences of each individual learner or team in the context of domain learning objectives* [1]. How an AIS might facilitate the accomplishment of domain learning objectives is illustrated in Fig. 2.

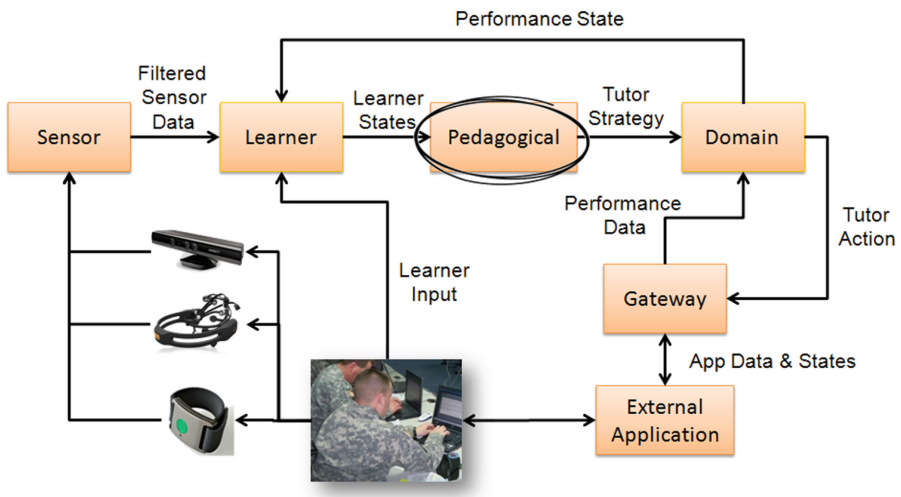


Fig. 2. Notional AIS architecture - tailoring is facilitated by the pedagogical model

Adaptive agents are one type of mechanism that facilitate adaptive instruction, and we have noted that adaptive agents take many forms and use many methods to adapt instruction. We emphasize that adaptive agents that act on an environment during instruction may or may not do so with the goal to enhance learning, and that not every adaptive instructional method is facilitated by an adaptive agent. To continue this discussion, we examine adaptation mechanisms in both the literature and in other papers within this volume.

2.1 Adaptation Mechanisms

In manually-adapted training systems, an instructor or simulation manager initiates adaptations when the conditions of the learner(s) warrant a change in the level of difficulty or the need for more or less support. In artificially-intelligent systems, adaptations are autonomously initiated/triggered by a set of conditions recognized by an AIS. Regardless of the nature of the instructor, the learning theory behind these adaptations is largely based upon Vygotsky’s Zone of Proximal Development (ZPD), a nexus within a learning experience where the difficulty level of the experience is perfectly balanced with the learner’s proficiency [12]. According to Vygotsky, this balance is necessary to keep the learner engaged during the learning experience (Fig. 3).

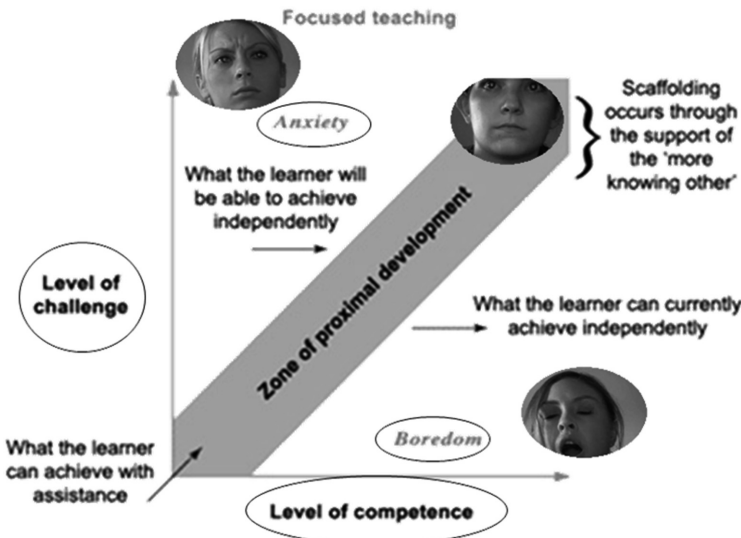


Fig. 3. Zone of proximal development (pictures courtesy of Graesser and D’Mello [13])

Jones and O’Grady (this volume) adopt a mixed-autonomy approach that allows all adaptations to be controlled on a spectrum from automated tuning to manual manipulation by human instructors. They discuss three types of adaptation: (1) tactics and techniques, e.g., adapting an attack or defense); (2) level of sophistication, e.g. making an attacker more or less aggressive, or limiting a defender’s awareness to focus training; and (3) personality parameters, e.g., tuning preferences of various types of agents in the ecosystem.

Freeman, Watz, and Bennett (this volume) discuss adaptive agents for adaptive tactical training. They emphasize the role of agents built using architectures capable of generating realistic and instructive behaviors. The authors specify the data required to

drive such agents and describe a testbed that delivers some of these data to accelerate development of agents and improve evaluation of agents' tactical smarts and agility. The adaptations in this case address tactical proficiency in battles against sophisticated human and machine adversaries, instructional efficacy, and maintenance and extension of agents to new training challenges. These adaptations span inferences fit to varied tactical situations, generation or selection of tactical actions to fit those inferences, use of different behavioral "modes" (e.g., the tactical preferences of different adversaries), fitting instructional actions to student needs, enabling agent evolution over time, and publishing metadata that enable systems to adaptively select and parameterize agents for new training tasks.

Warwick (this volume) advocates for adaptive agents that strike a balance between realism and tractability. The author argues that realism is subordinate to instructional objectives in training applications, and that in any specific instance, the level of realism required to meet those training objectives is an empirical question. For adaptive behaviors, he offers two general mechanisms. First, agents in training simulations (as opposed to simulation for other purposes) should have full visibility into "ground truth" and use that perspective to probe a learner's weaknesses or to create tactically challenging circumstances as teaching moments. Second, formalisms for encapsulating agent behaviors can be blended, so that, for instance, finite state machines can drive routine behaviors with richer cognitive representations injecting adaptive behaviors.

2.2 Defining Adaptive Instructional Agents

Now as we mix concepts of adaptive instruction and intelligent agents, we begin to formulate the characteristics of adaptive instructional agents. An adaptive agent can be thought of as modifying its behaviors to align with the current activity context, or adjusting its actions to learned user preferences. Adaptive instruction is generally speaking a learning interaction that adjusts to learner proficiencies, affective state, or other user characteristics [1, 6, 7, 14–16].

In examining the automated instructional decisions expected to be managed by intelligent agents, Sottolare distills them into three simple types: recommendations, strategies, and tactics [11]. Recommendations are relevant proposals that suggest possible next steps (e.g., next problem or lesson selection) and fit into what VanLehn [17] describes as the outer loop of the tutoring process, sometimes referred to as macro-adaptations. Strategies are plans for action by the AIS and tactics are actions executed by an AIS intelligent agent. Tactics are tied to VanLehn's inner loop [17] and are micro-adaptations used to guide step-by-step experiences (e.g., prompts, hints, or feedback during problem solving) in adaptive instruction.

As previously discussed, intelligent agents are autonomous entities which observe their environment through sensors and act upon their environment using actuators while directing its activity towards achieving goals [9]. In AISs, Baylor [18] identifies three primary functions or abilities for intelligent agents: (1) the ability to manage large amounts of data, (2) the ability to serve as a credible instructional and domain expert, and (3) the ability to modify the environment in response to learner performance. To this end, we add the requirement for the agent to be a learning agent, an entity that

makes more effective decisions with each experience. In a multi-agent AIS, we also add some additional capabilities prescribed by Nye and colleagues [19]:

- Specialization: agents can specialize in tasks they perform and can communicate their abilities using semantic messages and agent communication languages [20]
- Decentralization: a large number of agent services interacting across domains, platforms (e.g., desktop vs. mobile), and client and server-based software
- Customization: distinct agents that communicate using messages where each agent does not rely on the specific internal state of any other agent
- Data-driven enhancements: related to customization, where agents (and simpler services) use data to improve their performance against specified measures and goals

Finally, we acknowledge the requirement that adaptive instructional agents attenuate their “correctness” in order to incorporate instructionally meaningful and tactically plausible errors into the behaviors of synthetic teammates. The ability of an agent to commit deliberate errors at instructionally strategic moments can provide critical practice in team coordination skills [16].

The term “adaptive instructional agent” conveniently captures the convergence of adaptive agents that make recommendations or execute instructions (e.g., Google Assistant or Amazon Alexa) and applies their abilities to adaptive instruction. A key distinction here is that agents participating in instruction are, or could be, role-players in a scenario. Such agents thus play multiple functions: performing some role in the story; ensuring that instructional goals are achieved; providing real-time feedback; understanding and generating natural language. With this in mind, we define adaptive instructional agents as: *intelligent software-based entities that observe and act upon both the learner and the environment with the goal of enhancing learning with respect to allocated learning objectives*. It seems that we need and expect a lot from our adaptive instructional agents, but what impact does this have on the authoring process for AISs?

3 Implications for Authoring AIS

As we begin to evaluate the implications of a multi-agent architecture on the authoring process, we should begin with understanding the types of users directly interacting with authoring services to produce adaptive instruction. This following list was compiled from GIFT in its context as an evolving multi-agent architecture for authoring, delivering, managing, and evaluating adaptive instruction [19]:

- Instructors (Use Services): Deliver, modify, and possibly design GIFT courses
- Basic Course Designers (Configure Services): Modify or build course content with wizards/tutorials, including selecting or configuring a group of services verified to work together.
- Advanced Course Designers (Compose Services): Build advanced content and adaptive behavior, by selecting and configuring services to work together as a group

- Service/Agent Programmers (Make/Add Services): Code new services or agents used by GIFT
- Framework Providers (Combine Service Ecosystems): Maintain and interface other large frameworks, sensors, or external environments with GIFT through interfaces and gateways

Given these authoring roles, we envision authoring tools that link specified learning objectives with multi-agent services (e.g., assess learner performance or select optimal feedback). An AIS conceptual model including an ontology, a set of concepts and categories in a subject area or domain that shows their properties and the relations between them, is needed to identify variables for common services.

4 Adaptation Vectors

While not comprehensive, the following adaptation vectors have been identified in conjunction with the agent-based services associated with AISs:

- Learning objective services – authoring capabilities for automatically developing and tracking progress toward learning objectives
- Constructive behavioral services – authoring capabilities for developing, observing, and modifying adaptive entity behaviors for use in virtual environments stimulated by AISs; adaptations may be to provide more intelligent interactions and/or to provide more complex and difficult scenarios in response to learner performance
- Adaptive realism services – authoring capabilities that automatically scale the realism of the environment to align with the resolution required to perform the training task
- Learner modeling services – authoring capabilities for automatically populating individual learner or team models with variables that measure progress toward learning objectives and maintain learner states for learning, performance, retention, transfer of training, domain competency and/or proficiency
- Learner strategy services – authoring capabilities for automatically matching the difficulty of future experiences to align with learner proficiency (ala Vygotsky’s Zone of Proximal Development)
- Learner recommender services – automated capabilities to analyze organizational or career learning goals and recommend next steps for training, education, and job experiences based on gaps in knowledge or skills
- Dialogue-based and virtual human services – authoring capabilities for developing scripts, decision trees, and conditional models to drive dialogue based on available content

In addition to services, we might also seek to establish principles or tenets that govern the application and decision-making of adaptive instructional agents:

- Learning first - tactical realism (feasibility) of solutions are important, but may be sacrificed to create advantageous conditions for learning; an AIS might alter lesson sequences, level of difficulty, or interaction style, but may not necessarily apply such adaptations in ways that maintain the tactical believability of a scenario

- Long term learning first - adaptations must enhance learning outcomes, but long term learning goals are more important than short term learning goals; an AIS might forego short term learning successes to pursue long term learning goals
- Satisficing - timely solutions are important and timely solutions that are less optimal may be selected in lieu of less solutions; a timely good solution is better than a late optimal solution; an AIS might implement tactics that are less than optimal
- Longitudinal modeling - adaptations must be based not just on current states, but an examination of learner trends across longer periods of time; a comprehensive model of the learner and their progress toward learning goals is a better basis for decision-making than a snapshot

5 Conclusion and Recommendations

Designers of AISs and developers of the agents that inhabit them must reconcile a host of coincident and in some cases competing constructs that shape how instruction adapts to support the learner. In this paper we presented a notional set of principles to help instructional designers and agent developers implement adaptive instruction. Our specific focus is on the application of intelligent agents and their role in various tailoring methods, which we term *adaptation vectors*.

We distinguished agents for simulation (which act to advance scenario events but not necessarily to enhance learning) from agents for adaptive instruction, casting adaptation vectors as explicitly supporting both instructionally meaningful and doctrinally plausible actions by the AIS. We noted as well that “plausible” does not always equate to “correct” or “optimal”, since a legitimate adaptation to support learning objectives can steer an agent toward deliberately erroneous behavior.

As an initial framework, we proposed seven adaptation vectors within which AIS developers can define specific instances of adaptations: learning objectives; constructive behaviors; adaptive realism; learning modeling; learning strategy; learning recommender; and dialogue. As an early step toward granting agents the logic to select among these adaptation types, we further proposed four tenets that could be incorporated into adaptive agent logic: learning first; long-term learning first; satisficing; and longitudinal modeling.

In conclusion, we explored the literature to understand the design of adaptation vectors and provided recommendations for future AIS research and standards development. It should be noted that while not every adaptive instructional method is facilitated by an adaptive agent, many are and there is a trend toward learning agents. Our goal in setting forth this preliminary analysis is to encourage AIS developers to be deliberate in how adaptive agents influence instructional interactions. Further research is needed to validate these categories as effectively capturing distinctive and useful adaptation types.

References

1. Sottolare, R., Brawner, K.: Component interaction within the generalized intelligent framework for tutoring (GIFT) as a model for adaptive instructional system standards. In: The Adaptive Instructional System (AIS) Standards Workshop of the 14th International Conference of the Intelligent Tutoring Systems (ITS) Conference, Montreal, Quebec, Canada, June 2018
2. Oppermann, R.: *Adaptive User Support: Ergonomic Design of Manually and Automatically Adaptable Software*. Routledge, London (2017)
3. Oppermann, R., Rasher, R.: Adaptability and adaptivity in learning systems. *Knowl. Transf.* **2**, 173–179 (1997)
4. Sottolare, R.A., Brawner, K.W., Goldberg, B.S., Holden, H.K.: The generalized intelligent framework for tutoring (GIFT). In: Concept Paper Released as part of GIFT Software Documentation. U.S. Army Research Laboratory—Human Research & Engineering Directorate (ARL-HRED), Orlando (2012)
5. Sottolare, R., Brawner, K., Sinatra, A., Johnston, J.: An Updated Concept for a Generalized Intelligent Framework for Tutoring (GIFT). US Army Research Laboratory, Orlando (2017)
6. Anderson, J.R., Boyle, C.F., Reiser, B.J.: Intelligent tutoring systems. *Science* **228**(4698), 456–462 (1985)
7. Psotka, J., Massey, L.D., Mutter, S.A.: *Intelligent Tutoring Systems: Lessons Learned*. Lawrence Erlbaum Associates, Hillsdale (1988). ISBN 978-0-8058-0192-7
8. Baylor, A.: Beyond butlers: intelligent agents as mentors. *J. Educ. Comput. Res.* **22**(4), 373–382 (2000)
9. Russell, S., Norvig, P.: *Artificial Intelligent: A Modern Approach*. Pearson Education Ltd., Malaya (2003)
10. Niazi, M., Hussain, A.: Agent-based computing from multi-agent systems to agent-based models: a visual survey. *Scientometrics* **89**(2), 479–499 (2011)
11. Sottolare, R.A., Burke, C.S., Salas, E., Sinatra, A.M., Johnston, J.H., Gilbert, S.B.: Designing adaptive instruction for teams: a meta-analysis. *Int. J. Artif. Intell. Educ.* **28**(2), 225–264 (2018)
12. Vygotsky, L.S., Rieber, R.W.: *The Collected Works of LS Vygotsky: Volume 1: Problems of General Psychology, Including the Volume Thinking and Speech. Cognition and Language: A Series in Psycholinguistics, vol. 1*. Springer, Heidelberg (1987). <https://doi.org/10.1007/978-1-4613-1655-8>
13. Graesser, A.C., D’Mello, S.: Emotions during the learning of difficult material. In: Ross, B. (ed.) *Psychology of Learning and Motivation*, vol. 57, pp. 183–225. Academic Press, San Diego (2012)
14. Sottolare, R.: Considerations in the development of an ontology for a generalized intelligent framework for tutoring. In: *International Defense & Homeland Security Simulation Workshop in Proceedings of the I3M Conference*, Vienna, Austria, September 2012
15. Sottolare, R., Ragusa, C., Hoffman, M., Goldberg, B.: Characterizing an adaptive tutoring learning effect chain for individual and team tutoring. In: *Proceedings of the Interservice/Industry Training Simulation & Education Conference*, Orlando, Florida (2013)
16. Bell, B., McNamara, J.: How to avoid using stupid agents to train intelligent people. In: *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (IITSEC)*, December 2005
17. VanLehn, K.: The behavior of tutoring systems. *Int. J. Artif. Intell. Educ.* **16**(3), 227–265 (2006)

18. Baylor, A.: Intelligent agents as cognitive tools for education. *Educ. Technol.* **1**, 36–40 (1999)
19. Nye, B.D., Auerbach, D., Mehta, T.R., Hartholt, A.: Building a backbone for multi-agent tutoring in GIFT (Work in progress). In: *Proceedings of the 5th Annual Generalized Intelligent Framework for Tutoring (GIFT) Users Symposium (GIFTSym5)*, 17 July 2017, p. 23 (2017)
20. Bellifemine, F., Caire, G., Poggi, A., Rimassa, G.: JADE: a software framework for developing multi-agent applications. *Inf. Softw. Technol.* **50**(1), 10–21 (2008)