

A Conceptual Framework for Human–AI Hybrid Adaptivity in Education

Kenneth Holstein¹, Vincent Aleven¹, and Nikol Rummel²

¹ Carnegie Mellon University, Pittsburgh PA 15213, USA

² Ruhr-Universität Bochum, 150, 44801 Bochum, Germany

{kjholste, aleven}@cs.cmu.edu, nikol.rummel@rub.de

Abstract. Educational AI (AIEd) systems are increasingly designed and evaluated with an awareness of the hybrid nature of adaptivity in real-world educational settings. In practice, beyond being a property of AIEd systems alone, adaptivity is often *jointly enacted* by AI systems and human facilitators (e.g., teachers or peers). Despite much recent research activity, theoretical and conceptual guidance for the design of such human–AI systems remains limited. In this paper we explore how adaptivity may be shared across AIEd systems and the various human stakeholders who work with them. Based on a comparison of prior frameworks, which tend to examine adaptivity in AIEd systems or human coaches separately, we first synthesize a set of dimensions general enough to capture *human–AI hybrid adaptivity*. Using these dimensions, we then present a conceptual framework to map distinct ways in which humans and AIEd systems can augment each other’s abilities. Through examples, we illustrate how this framework can be used to characterize prior work and envision new possibilities for human–AI hybrid approaches in education.

Keywords: Adaptivity, Human–AI Hybrid, Orchestration, Collaboration, Framework

1 Introduction

Moving beyond a focus on adaptivity as a property of AIEd systems alone, AIEd research increasingly acknowledges that, in practice, adaptive learning experiences may be *jointly enacted* by AI and human facilitators (e.g., [7; 15; 24; 30; 47; 58; 70]). For instance, recent work indicates that in K-12 classrooms using AI tutoring software, the sequence of educational activities students receive is often driven by a combination of AI-based activity selection and the dynamic decision-making of classroom teachers (who may selectively override algorithmic recommendations) [53]. Other work has explored the nature and impacts of human–human interactions during AI-supported class sessions, finding that these interactions can play critical roles in mediating AIEd technologies’ effectiveness [25; 26; 30; 41; 48; 70]. Building upon such observations, a number of recent projects have begun to explore how AIEd systems might more effectively *work together* with human facilitators, to amplify their abilities and leverage their complementary strengths [18; 24; 26; 42; 47; 66; 70].

As the AIEd community increasingly turns its attention to *human–AI hybrid* approaches for education, some conceptual guidance may be helpful in navigating this broad design space and in differentiating between fundamentally different kinds of hybrid approaches. Different configurations of AIEd systems and humans, designed to

integrate human and AI abilities in different ways, may yield very different outcomes (e.g., [26; 55; 66; 70]). In this paper, we begin to map the diverse ways in which adaptivity may be shared among humans and AIED systems, to aid the community in (1) organizing *prior work* through the lens of human–AI hybrid adaptivity, and (2) envisioning *new possibilities* for human–AI hybrid approaches in education. To this end, we present a conceptual framework for human–AI hybrid adaptivity in education. Drawing upon multiple existing frameworks for adaptive support—here defined broadly as support that is responsive to unfolding learning situations in pursuit of educational goals—we begin by synthesizing a set of dimensions general enough to capture human–AI adaptivity (Section 2). Using these dimensions, we then introduce distinct ways in which humans and AI might augment each other’s abilities, illustrating the framework’s utility via examples of new directions it surfaces (Section 3).

2 Framing “Adaptivity”: Synthesizing Existing Frameworks

Several frameworks have been developed to characterize adaptivity in education. In this paper, we build upon a small set of prior frameworks [3; 21; 50; 51; 56; 57; 61; 65] to inform our thoughts about what a more encompassing framework should include. In selecting this set we aimed to consider influential work across multiple research areas, including AIED [3; 21; 50; 65], computer-supported collaborative learning (CSCL) [56; 61; 65], teacher cognition [57], and classroom orchestration [51]. We searched broadly for theoretically oriented articles that focus on characterizing *adaptive instructional behavior*. While the resulting selection of prior frameworks is not intended to be exhaustive, this set presents several interesting contrasts and overlaps.

Each of the frameworks considered offers a lens to examine particular aspects of adaptive learning systems, while abstracting over others. As discussed below, some frameworks, such as the Adaptivity Grid [3] and Plass’s framework [50] provide high resolution lenses to analyze *what* an adaptive system might respond to and *when* an adaptive system might respond, but do not, for example, offer explicit language for describing *how* an adaptive system might respond (see **Action space** below). Meanwhile, other frameworks focus much of their resolution towards characterizing the design space for instructional support actions. For example, VanLehn [65] and Rummel [56] offer ways of characterizing *how* and *when* a system might respond, yet do not offer language for *what* to respond to (see **Perceptual capabilities** below). One possible reason for these differences is that different frameworks have tended to focus on different kinds of adaptive learning systems. A related possibility is that because different frameworks are grounded in different research literatures (e.g., CSCL versus AIED [65]) they are heavily influenced by the state of the empirical literature within each community. For example, the Adaptivity Grid [3] offered finer-grained distinctions in areas where there was much existing empirical work at the time of writing, but offered coarser-grained distinctions where less prior work existed.

In the remainder of this paper, we adopt a broad framing of adaptivity in terms of *perception-action cycles* [11; 44; 62; 65] enacted by decision-making *agents* or systems of agents (e.g., AI, students, and teachers) [56], in service of specified educational *goals* [56; 65]. Building from prior frameworks, in this section we provide a set of dimensions that are general enough to encompass prior frameworks, while also providing language

rich enough to characterize a broad possibility space for human–AI hybrid adaptivity. Whereas prior frameworks focus on providing partial views of agents’ adaptive behavior, as discussed above, our dimensions draw from multiple frameworks to provide a more encompassing perspective (cf. [43]). At the same time, we abstract over dimensions from these prior frameworks in the interest of generalizing across a broad range of instructional systems and contexts. For instance, six of the dimensions proposed in [56] are collapsed into the *Actions* dimension below, given that all of these dimensions capture properties of instructional support actions in CSCL.

Goals and Targets: Adaptive instruction presupposes educational goals, or outcomes that the adaptive behavior is intended to bring about (which may vary by student or group and may change over time). For example, some AIED systems may be designed to adapt instruction with the ultimate goal of improving student learning outcomes within a domain, whereas others may adapt with the goal of helping students become better self-regulated learners or collaborators. Notably, only some prior frameworks for adaptive instruction provide vocabulary to describe the end goal(s) of the adaptivity. Rummel [56] explicitly names *goals* as the first dimension that needs to be defined upfront of designing any support. Both Rummel [56] and VanLehn [65] further distinguish between the ultimate *goals* of the support (e.g., the kind of change the adaptivity is intended to produce in students), and the immediate *targets* of the support (e.g., whether the support targets cognitive versus metacognitive knowledge).

Perceptual capabilities: Decision-making agents can adapt to unfolding learning situations only to the extent that they can perceive (i.e., sense and interpret [11; 20]) and represent these situations. An agent’s ability to perceive particular variables of a learning situation defines what it can potentially *adapt to*. In addition to variables that are directly observable, this may also include ones that the agent is able to infer from observable attributes (e.g., inferring a student’s or teacher’s current knowledge from patterns in their recent behavior). In an *Intelligent Tutoring System (ITS)*, the system’s perceptual capabilities are defined by its student modeling capabilities, which may include unobservable, inferred constructs such as “help avoidance” or “frustration” [13; 21; 29]. A *human teacher’s* perceptual capabilities can be understood as the range of phenomena the teacher is capable of sensing and inferring about a learning situation. In realistic contexts, this may depend on factors such as the teacher’s current attentional load [51; 52], as well as the teacher’s skill in noticing instructionally relevant events and drawing correct inferences based on potentially limited observations [51; 57; 59]. As noted above, some, but not all prior frameworks included explicit language to characterize an adaptive agent’s perceptual capabilities. The Adaptivity Grid [3] categorized previously published empirical evaluations of adaptive learning technologies, in part, based on whether they adapt instruction based on perceptions of students’ *prior knowledge & knowledge growth*, their *path through an activity*, their *affective & motivational states*, their *SRL strategies, metacognition, & effort*, or based on a notion of *learning styles*. Similarly, Plass (2016) categorized adaptive learning technologies based on whether they adapt instruction based on perceptions of *affective, cognitive, motivational, or socio-cultural* variables [50].

Action space: An agent’s ability to adapt instruction is also delimited by the set of responses or instructional moves it has at its disposal [56; 57; 61; 62; 65]. For instance, an ITS or a human tutor might try to adapt the kinds of help they provide to a student in their class based on their perceptions of the student’s current knowledge state.

However, the tutor’s ability to adapt will be limited by the instructional moves they currently have in their repertoire (e.g., providing correctness feedback, presenting a worked example, or prompting a self-explanation). Some, but not all, of the frameworks we reviewed included dimensions to characterize an agent’s action space. Soller [61] and VanLehn [65] distinguish between actions that *mirror* an agent’s perceptions back to students or human facilitators, actions that present an agent’s *assessments* of what it perceives, and *coaching* actions (e.g., providing advice). Rummel [56] presents multiple related dimensions classifying instructional support actions, for instance the *directivity* of an action (i.e., whether and to what extent the action presents explicit guidance). In addition, VanLehn [65] and Rummel [56] both characterize instructional actions in terms of their *recipient* or *addressee* (e.g., whether a system presents information to a student, a group of students, or an instructor), and Rummel further specifies whether a student (or group of students) is the direct target of an action, or whether the action is *mediated* through other actors in the learning environment (e.g., where an adaptive system suggests that a teacher or peer tutor help a given student).

Decision policies: An agent’s adaptive behavior can be understood in terms of *decision policies*: sets of rules that map (in a potentially non-deterministic manner) from perceived learning situations or states to particular actions that the agent will take in response [62]. For example, an agent might adaptively respond to detected student frustration by acknowledging or mirroring the student’s frustration [21; 50; 65]. However, many alternative decision policies exist. The system might instead respond to detected frustration by selecting alternative activities for the student to work on, or by asking the student whether the system should alert their teacher/peers that they need help [28]. Prior frameworks do not typically provide explicit dimensions to categorize “types” of decision policies (e.g., “responding to affect with affective responses” or “mastery learning based activity selection policies”), although such categorizations often appear in practice when empirically comparing different forms of adaptivity.

Granularity and Timing: Finally, many prior frameworks provide dimensions dedicated to describing *when* a system adapts instruction (e.g., [3; 50; 51; 56; 65]). That is, the frequency or granularity at which the perception-action cycle is enacted. This may occur, for instance, once per *task* or per *step* of a task [3; 56; 65], once per *turn* in a conversation [56], or even once per *design iteration* (when considering systems that are iteratively improved based on data) [3]. Plass [50], Prieto [51], and Rummel [56] also distinguish the *timing* of the adaptation; e.g., whether the adaptation occurs *prior* to the instructional activity, *in the midst* of the activity, or *afterwards* [50; 51; 56].

Many frameworks for adaptive instructional support have been developed, with each offering a lens to examine particular aspects and particular kinds of adaptive learning systems. The set of high-level dimensions presented in this section are intended to capture essential components of adaptive learning systems, informed by a comparison across frameworks (cf. [43]). In the next section, we use these dimensions to explore distinct ways for adaptivity to be *shared* across humans and machines.

3 A Conceptual Framework for Human–AI Hybrid Adaptivity

In the following we present a conceptual framework for human–AI hybrid adaptivity in education, examining the same set of basic components (goals/targets, perception,

action, decision policies, and granularity/timing) while broadening our focus. We use this framework both to characterize prior work and to envision new possibilities, based upon distinct ways in which humans and AIEd systems might augment one another: (1) *Goal Augmentation*, (2) *Perceptual Augmentation*, (3) *Action Augmentation*, and (4) *Decision Augmentation*. Within each category, possibilities exist both for *augmenting performance* (in which humans and AI systems, assumed to have complementary strengths and weaknesses, augment one another’s abilities at “runtime”, but without necessarily producing lasting changes in behavior) and for *co-learning* (in which humans and AI systems help one another improve over time). Finally, we discuss how the *Granularity and Timing* of adaptivity might be understood in human–AI systems.

3.1 Goal Augmentation: Informing Each Other’s Instructional Goals

A key way for humans and AIEd systems to support one another is by influencing each other’s goals. To a large extent, AIEd technologies encode the assumptions and goals of those who design and develop them—whether explicitly, via objective functions that a system’s adaptive policies optimize towards, or implicitly, through design decisions that promote certain goals over others. However, the goals baked into an AIEd system may not always align with those of humans in real-world educational contexts [24; 46; 53]. For example, ITSs used in K-12 school contexts often implement mastery-based activity selection policies, allowing each student to progress through the curriculum at their own pace. Yet prior work suggests that teachers often struggle to balance their desire to implement such personalized classrooms with external pressure to keep classes “on schedule”. In practice, teachers often opt to manually push students forward in the curriculum if they are slower to master certain skills [24; 53], sometimes even if they are aware that doing so may harm students’ learning [24; 28]. As of yet, little work in AIEd has explored the design of supports for goal augmentation.

AIEd informing human goals. It may not *always* be desirable for AIEd systems to adapt to human facilitators’ instructional goals. For instance, in some cases, teachers’ or peer tutors’ goals may be fundamentally at odds with known instructional best practices. Future systems could play an important role in helping humans productively reflect upon their goals, helping them refine these goals or consider alternatives [4; 19].

Humans informing AIEd goals. Human facilitators may hold critical, on-the-ground knowledge about their instructional contexts and personal goals, to which AIEd systems would not typically be privy. Building upon the above example, ITSs might be even more effective in classroom contexts if designed to accept teachers’ input regarding the goals they should be optimizing towards. By enabling teachers to help shape the system’s goals, the system could in turn help teachers more effectively navigate trade-offs between competing goals (e.g., by supporting teachers in deciding when to push students ahead in the curriculum, while causing minimal harm to their learning [28]).

3.2 Perceptual Augmentation: Leveraging Complementarity in Perception

A second way for AIEd systems and humans to augment one another is by enhancing each other’s abilities to perceive instructionally relevant information, or *opportunities for action*. This may take the form of (1) extending what the other is able to *sense* (i.e., what information is made available to them, prior to further interpretation [11; 20]); (2)

guiding how the other distributes their *attention*; or (3) guiding how the other *interprets* incoming information. Each of these broad possibilities is discussed in turn, below.

Augmenting sensing and attention.

AIED systems can be designed to extend what humans are able to *sense* and *notice* about learners, learning, or their own teaching, or from the other direction, to help humans augment what AIED systems sense and notice. Thus far, more work in AIED has focused on supporting *AI*→*human* than *human*→*AI* augmentation in this area.

AIED augmenting human sensing and attention. A number of AIED systems have been developed to help human facilitators sense information to which the AI would otherwise have unique access (e.g., [2; 5; 26; 38; 40; 55; 69; 70]). Prior work has focused on augmenting what learners and peer tutors are able to sense and notice about a learning situation. For example, the *Adaptive Peer Tutoring Assistant (APTA)* supports peer tutors in recognizing opportunities for effective intervention, in the context of ongoing peer tutoring [70]. In the context of self-regulated learning with an AI tutor, the *Help Tutor* supports students in monitoring their own help-seeking behavior, and in noticing cases where they may be using the software’s help functions in maladaptive ways [2]. More recently, several projects have focused on designing ways to keep human teachers in the loop in AI-supported classrooms (e.g., [27; 40; 47; 68]). For example, the *Lumilo* teacher smartglasses are designed to direct teachers’ attention, during a class session, to situations that an AI tutor may be poorly suited to handle on its own, or which require a teacher’s further assessment [26; 27]. In each of the above examples, there is potential for future AIED systems not only to augment human facilitators’ abilities in-the-moment, but also to help humans *learn to notice* relevant features of a learning situation even when in-the-moment support is unavailable [2; 19; 59; 70].

Humans augmenting AIED sensing and attention. From the other side, humans may have relevant on-the-ground knowledge to which AIED systems are likely to be blind. AIED systems may be designed so that humans can help them perceive such information. For instance, future systems might be designed to allow teachers and parents to update individual student models with relevant information about the student’s broader context; e.g., whether the student is currently facing at-home difficulties that may impact their performance (cf. [9]). Similarly, an AIED system might be designed to periodically poll students regarding their subjective *feeling of knowing* particular skills that are targeted by the instruction [12; 36]. In addition to having humans input information directly, some research has begun to explore approaches in which humans teach AIED systems, via demonstration, to perceive instructionally relevant features to which they should attend in the future (e.g., [35]).

Augmenting interpretation.

AIED systems can also be designed to support humans in *interpreting* and *drawing inferences* from what they notice, or to assist humans in shaping or mediating AIED systems’ interpretations of the events they are able to sense.

AIED augmenting human interpretation. Beyond extending human sensing capacities, AIED systems may also support human facilitators in productively interpreting and reflecting upon the information available to them. Whereas some technologies are designed to present information to humans with low-level, minimally pre-interpreted data (e.g., “number of help requests”) [4; 5; 16], several of the AIED systems discussed

above, including *APTA*, the *Help Tutor*, and *Lumilo* rely upon advanced student modeling techniques (e.g., automated detectors of “help abuse” or “help avoidance” behaviors). Thus, beyond augmenting human sensing and attention, these systems perform a considerable amount of pre-interpretation on behalf of human facilitators or learners [5]. Emerging lines of research are beginning to explore the design of interfaces that can more actively guide humans towards particular interpretations of learning data (e.g., [17]) or interfaces that can scaffold humans in more productive forms of reflection (e.g., [19]). However, it remains an open question for future research how best to productively guide human interpretation, while still leveraging (rather than diminishing) humans’ unique inferential capacities [5; 14; 17; 23; 33].

Humans augmenting AIEd interpretation. Future AIEd systems may be designed to support human facilitators in *detecting* cases where the AI misinterprets learning data (e.g., by misclassifying patterns in collaborating groups’ behaviors) and to *provide corresponding feedback* in order to shape these interpretations in more meaningful directions. As of yet, the question of how AIEd systems can be designed to effectively elicit and learn from such feedback remains underexplored (cf. [9; 10; 14; 23; 34; 54]).

3.3 Action Augmentation: Leveraging Complementarity in Action Spaces

A third way in which AIEd systems and humans can work together to support more adaptive instruction is by augmenting and extending the other’s capacities for instructional *action*. In particular, AIEd systems and humans can (1) enhance each other’s *ability to perform* particular kinds of instructional actions, and relatedly *expand the range of actions available* to each; and (2) enhance each other’s *scalability* and *capacity* for action. Each of these broad possibilities is discussed below.

Enhancing ability and expanding the availability of actions.

Many open research and design opportunities exist for human–AI systems that augment and expand each other’s action spaces. Just a few examples are presented below.

AIEd augmenting human actions. AIEd systems may be designed to support human facilitators in providing *more effective* help. For example, while a human coach works with a student, a future AIEd system might follow along with what the coach is doing, and adaptively present educational resources (e.g., relevant readings, videos, or practice materials) that support their current goals [28; 70]. Alternatively, a system may respond during or after human coaching by adaptively providing feedback on the quality of the instruction (e.g., the clarity of a particular explanation the coach provides), to help the coach adjust and improve over time (cf. [19; 28; 70]).

Humans augmenting AIEd actions. Humans can also augment the set of instructional moves available to an AIEd system by either *customizing* or creating *new actions* for the system. For example, AIEd systems may be designed to adaptively deliver hints written by peers or instructors (cf. [22; 28; 72]). Authoring tools have been developed to support non-programmer authoring, but further research is needed to support easy authoring in everyday educational settings (e.g., by teachers or students) [1; 29; 37; 39].

Enhancing scalability and capacity.

Much prior research in AIEd has focused on augmenting human scalability, whereas relatively less research has targeted the reverse direction. However, many open questions remain in each direction, which emerging work is beginning to tackle.

AIEd augmenting human scalability. AIEd systems have often been promoted as “scaling up” some of the benefits of one-on-one tutoring, effectively providing each student with their own, personal AI tutor [6; 31; 58; 64]. In doing so, AIEd systems can serve as *teachers’ aides* [24; 73], helping human coaches or teachers personalize instruction beyond what might otherwise be feasible, while also freeing up humans’ limited time and attention for other activities (e.g., providing socio-emotional support or coaching for students most in need) [22; 24; 58; 73]. Thus, one way in which AIEd systems can augment human scalability and capacity is through selective *delegation* [27]. Some research has begun to explore the design of AIEd systems that adaptively, dynamically delegate instructional roles between AI systems, teachers, and peers, based upon an awareness of trade-offs between the *instructional ability* and *capacity* of each [28; 47; 49; 63]. A second emerging way for AIEd systems to help human facilitators scale their efforts is by supporting them in *teaching the AI tutor* (as discussed below), transferring their unique expertise and pedagogical preferences into a system that can reach more students than they themselves can [37, 39, 60, 74].

Humans augmenting AIEd scalability. It can also be useful to consider the ways in which human facilitators can (and in practice, often do) support AIEd systems in scaling. Increased scalability risks reducing a system’s fit with particular educational contexts, as system developers design solutions to fit constraints of multiple contexts simultaneously [32; 45; 46]. On-the-ground facilitators may support AIEd systems in scaling to diverse contexts by adapting the way these systems are implemented in use to the needs of their local contexts (e.g., [24; 30; 46; 58]). For example, when classroom teachers use AIEd systems that are poorly aligned with their school’s existing curriculum, they may selectively assign particular modules to students, overriding the systems’ built in sequencing algorithms in the interest of providing better aligned learning experiences [23; 45]. Future AIEd systems may be explicitly designed to facilitate such *adaptability* (e.g., local customizations and overrides) [16], improving their chances for adoption across varied contexts of use [22; 24; 28; 45].

3.4 Decision Augmentation: Leveraging Complementarity in Decision-making

Beyond informing each other’s goals or augmenting each other’s capacities for perception and action, a fourth major way in which AIEd systems and humans can work together is by helping each other make *more effective pedagogical decisions* (i.e., helping each other more effectively link between perception and action). Prior work has explored forms of both *AI→human* and *human→AI* decision augmentation. However, much additional research is needed in order to fully realize the visions of AIEd systems as, for instance, effective *decision support* and *professional development* tools [5; 19; 23; 24; 66] and as *teachable machines* [37; 39; 60; 74].

AIEd augmenting human decision-making. In addition to providing instruction to students directly, AIEd systems may be designed as decision support for human facilitators, helping humans take more effective instructional actions in particular learning situations [2; 26; 27; 66; 67; 70]. To an extent, all forms of human augmentation discussed thus far can function as forms of decision support. Indeed, decision support is often conceptualized as a continuous spectrum rather than a binary design choice [5; 56; 65; 71]. For instance, perceptual augmentation may enhance decision-making by directing humans’ attention towards learning phenomena that

require their further assessment or action [5]. However, AIED systems may also be designed to support human decision-making more directly and explicitly. For example, an AIED system might automatically suggest effective ways for a human facilitator to help a group of students, in the moment, based on its perceptions of the students' and/or facilitator's current states (effectively functioning as hints or bug messages, targeted for a human in an instructional role rather than a learner; see [27; 28; 68; 70]). With knowledge of a facilitator's instructional goals, future AIED systems might help the facilitator make more informed trade-offs between competing goals or nudge them away from practices that are at odds with their goals [28]. Such systems could function not only as decision support, but also as professional development, helping humans improve over time, potentially even in the absence of such support [19; 27; 28; 70].

Humans augmenting AIED decision-making. AIED systems may also be designed to help human facilitators *mediate* or *shape* these systems' instructional decision-making. Mediation may occur in practice where a facilitator such as a teacher overrides a decision made by an AIED system (e.g., by selecting an alternative activity for a student to work on, or an alternative group for a student to collaborate with, rather than ones selected by the system) [5; 24; 47; 49]. As discussed under *Goal Augmentation* above, such overriding behavior occurs regularly in K-12 classroom contexts; as noted, although teachers' overrides can often be seen as adaptive, they can also be maladaptive when they detract from (some) goals for the instruction. In addition to mediating AIED systems' decision-making, humans might also help systems learn more effective policies or ones better suited to their particular educational contexts [23]. Recent work on *machine teaching* for AIED suggests promise for approaches in which humans *teach the AI to teach* through feedback and demonstrations [37; 39]. However, further research is needed to develop interaction paradigms for machine teaching that are fast and intuitive enough for everyday use in educational settings [60; 74].

3.5 Granularity and Timing in Human–AI Systems

Finally, we briefly discuss how *granularity* and *timing* might be understood in human–AI systems. When AIED systems and humans work together, they may each adapt instruction at different grain sizes. For instance, in classrooms using step-based tutoring software, teachers may provide *substep* feedback on-the-spot (i.e., feedback on a step while the student is, from the system's perspective, still in the midst of completing the step) [23; 27; 30]. While an AIED system waits for the student to submit their input, a human facilitator might perceive an opportunity to intervene within a long pause in student typing. The *timing* of adaptation may also vary across humans and machines. For instance, Aleven et al's "design loop adaptivity" [3] can be viewed as involving a form of shared adaptivity in which human facilitators or instructional designers repeatedly adapt an AIED system's design (informed by educational data and/or their own observations) *before or after* an instructional activity, while the AIED system in turn takes care of adapting to learning situations *during* the activity.

4 Conclusions

AIED systems are increasingly designed and evaluated with an awareness of the shared nature of adaptivity in real-world educational settings. Despite much recent research

into human–AI hybrid approaches for education, theoretical and conceptual guidance in this area remains limited. Whereas prior frameworks have tended to examine adaptivity in AIEd systems or human coaches separately, in this paper we have explored how adaptivity may be shared across AIEd systems and the various human stakeholders who work with them.

Based on a comparison and synthesis of prior frameworks, we have presented a generalized set of dimensions, with the goal of capturing essential components of adaptive instructional behavior (cf. [43]). Using these dimensions, we have introduced a conceptual framework for *human–AI hybrid adaptivity* in education, suggesting distinct ways in which AIEd systems and human facilitators might augment one another. Throughout the previous section, we have presented several examples to illustrate how this framework can be used both to characterize *prior work* and to surface *new possibilities* and *open questions* for human–AI hybrid approaches in education.

We view the current framework as a step towards the development of richer theory for human–AI hybrid adaptivity in education, and for human–AI hybrid approaches more broadly. As an empirical and design science, AIEd needs theory to productively guide hypothesis generation, prediction, understanding, and design. Theory can help researchers adopt common concepts and vocabulary, which may in turn accelerate communication and innovation. Theory can shape—for better or worse—how researchers and designers see the world, how they make sense of their observations, and what alternatives they are able to envision. The current framework should be viewed as a starting point, not a finished product. We invite others in the community will challenge this framework and expand upon it.

The design space for human–AI hybrid approaches in education is large and combinatorial: almost any real case will involve *combinations* of the categories of human–AI adaptivity specified in this framework (e.g., an AIEd system might augment human decision-making via a human-augmented perceptual model). It is our hope that the present work will help to guide future research and design, assisting others in navigating this broad design space, in formulating more useful hypotheses, and in differentiating among fundamentally different kinds of human–AI hybrid approaches.

Acknowledgements

This work was supported by NSF Grant #1822861 and IES Grant R305A180301. Any opinions expressed in this article are those of the authors and do not represent views of the NSF or IES.

References

1. Aleven, V., McLaren, B. M., Sewall, J., Van Velsen, M., Popescu, O., Demi, S., Ringenberg, M., & Koedinger, K. R. (2016). Example-tracing tutors: Intelligent tutor development for non-programmers. *International Journal of Artificial Intelligence in Education*, 26(1), 224-269.
2. Aleven, V., Roll, I., McLaren, B. M., & Koedinger, K. R. (2016). Help helps, but only so much: Research on help seeking with intelligent tutoring systems. *International Journal of Artificial Intelligence in Education*, 26(1), 205-223.

3. Aleven, V., McLaughlin, E. A., Glenn, R. A., & Koedinger, K. R. (2016). Instruction based on adaptive learning technologies. *Handbook of research on learning and instruction*, 522-560.
4. An, P., Bakker, S., Ordanovski, S., Taconis, R., Paffen, C. L., & Eggen, B. (2019). Unobtrusively enhancing reflection-in-action of teachers through spatially distributed ambient information. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1-14).
5. An, P., Holstein, K., d'Anjou, B., Eggen, B., & Bakker, S. (2020). The TA Framework: Designing Real-time Teaching Augmentation for K-12 Classrooms. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI'20)*.
6. Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences*, 4(2), 167-207.
7. Baker, R. S. (2016). Stupid tutoring systems, intelligent humans. *International Journal of Artificial Intelligence in Education*, 26(2), 600-614.
8. Bransford, J. D., Brown, A. L., & Cocking, R. R. (2000). *How people learn* (Vol. 11). Washington, DC: National academy press.
9. Bull, S., & Kay, J. (2016). SMILI©: A framework for interfaces to learning data in open learner models, learning analytics and related fields. *International Journal of Artificial Intelligence in Education*, 26(1), 293-331.
10. Chen, N. C., Suh, J., Verwey, J., Ramos, G., Drucker, S., & Simard, P. (2018). AnchorViz: Facilitating classifier error discovery through interactive semantic data exploration. In *23rd International Conference on Intelligent User Interfaces* (pp. 269-280).
11. Creem-Regehr, S. H., & Kunz, B. R. (2010). Perception and action. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(6), 800-810.
12. Cromley, J., Azevedo, R., & Olson, E. (2005). Self-Regulation of Learning with Multiple Representations in Hypermedia. In *Proceedings of the International Conference on Artificial Intelligence in Education* (pp. 184-191). Springer.
13. Desmarais, M. C., & Baker, R. S. (2012). A review of recent advances in learner and skill modeling in intelligent learning environments. *User Modeling and User-Adapted Interaction*, 22(1-2), 9-38.
14. De-Arteaga, M., Fogliato, R., & Chouldechova, A. (2020). A Case for Humans-in-the-Loop: Decisions in the Presence of Erroneous Algorithmic Scores. *arXiv preprint arXiv:2002.08035*.
15. Dillenbourg, P. (2016). The evolution of research on digital education. *International Journal of Artificial Intelligence in Education*, 26(2), 544-560.
16. Dillenbourg, P., Nussbaum, M., Dimitriadis, Y., & Roschelle, J. (2013). Design for classroom orchestration. *Computers & Education*, 69(0), 485-492.
17. Echeverria, V., Martinez-Maldonado, R., Shum, S. B., Chiluiza, K., Granda, R., & Conati, C. (2018). Exploratory versus Explanatory Visual Learning Analytics: Driving Teachers' Attention through Educational Data Storytelling. *Journal of Learning Analytics*, 5(3), 72-97.
18. Fancsal, S. E., Yudelson, M. V., Berman, S. R., & Ritter, S. (2018). Intelligent Instructional Hand Offs. *International Educational Data Mining Society*.
19. Gerritsen, D., Zimmerman, J., & Ogan, A. (2018). Towards a framework for smart classrooms that teach instructors to teach. *International Society of the Learning Sciences, [ISLS]*.
20. Goldstein, E. B., & Brockmole, J. (2016). *Sensation and perception*. Cengage Learning.
21. Harley, J. M., Lajoie, S. P., Frasson, C., & Hall, N. C. (2017). Developing emotion-aware, advanced learning technologies: A taxonomy of approaches and features. *International Journal of Artificial Intelligence in Education*, 27(2), 268-297.
22. Heffernan, N. T., & Heffernan, C. L. (2014). The ASSISTments ecosystem: Building a platform that brings scientists and teachers together for minimally invasive research on human learning and teaching. *International Journal of Artificial Intelligence in Education*, 24(4), 470-497.

23. Holstein, K. (2019). Designing Real-time Teacher Augmentation to Combine Strengths of Human and AI Instruction. *Unpublished doctoral dissertation, Carnegie Mellon University*.
24. Holstein, K., McLaren, B. M., & Aleven, V. (2017). Intelligent tutors as teachers' aides: Exploring teacher needs for real-time analytics in blended classrooms. In *Proceedings of the Seventh International Learning Analytics & Knowledge Conference* (pp. 257-266).
25. Holstein, K., McLaren, B. M., & Aleven, V. (2017). SPACLE: Investigating learning across virtual and physical spaces using spatial replays. In *Proceedings of the Seventh International Learning Analytics & Knowledge Conference* (pp. 358-367).
26. Holstein, K., McLaren, B. M., & Aleven, V. (2018). Student learning benefits of a mixed-reality teacher awareness tool in AI-enhanced classrooms. In *International Conference on Artificial Intelligence in Education* (pp. 154-168). Springer, Cham.
27. Holstein, K., McLaren, B. M., & Aleven, V. (2019). Co-designing a real-time classroom orchestration tool to support teacher-AI complementarity. *Journal of Learning Analytics*, 6(2), 27-52.
28. Holstein, K., McLaren, B. M., & Aleven, V. (2019). Designing for complementarity: Teacher and student needs for orchestration support in AI-enhanced classrooms. In *International Conference on Artificial Intelligence in Education* (pp. 157-171). Springer, Cham.
29. Holstein, K., Yu, Z., Sewall, J., Popescu, O., McLaren, B. M., & Aleven, V. (2018). Opening up an intelligent tutoring system development environment for extensible student modeling. In *International Conference on Artificial Intelligence in Education* (pp. 169-183). Springer, Cham.
30. Kessler, A., Boston, M., & Stein, M. K. (2019). Exploring how teachers support students' mathematical learning in computer-directed learning environments. *Information and Learning Sciences*.
31. Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent Tutoring Goes To School in the Big City. *International Journal of Artificial Intelligence in Education*, 8, 30-43.
32. Kulkarni, C. (2019). Design Perspectives of Learning at Scale: Scaling Efficiency and Empowerment. In *Proceedings of the Sixth (2019) ACM Conference on Learning@ Scale* (pp. 1-11).
33. Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40.
34. Lakkaraju, H., Kamar, E., Caruana, R., & Horvitz, E. (2017). Identifying unknown unknowns in the open world: Representations and policies for guided exploration. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*.
35. Lee, M. H., Runde, J., Jibril, W., Wang, Z., & Brunskill, E. (2015). Learning the features used to decide how to teach. In *Proceedings of the Second (2015) ACM Conference on Learning@ Scale* (pp. 421-424).
36. Long, Y., & Aleven, V. (2011). Students' understanding of their student model. In *Proceedings of the International Conference on Artificial Intelligence in Education* (pp. 179-186). Springer, Berlin, Heidelberg.
37. Maclellan, C. J., Harpstead, E., Patel, R., & Koedinger, K. R. (2016). The Apprentice Learner Architecture: Closing the Loop between Learning Theory and Educational Data. *International Educational Data Mining Society*.
38. Martinez-Maldonado, R., Clayphan, A., Yacef, K., & Kay, J. (2014). MTFeedback: Providing notifications to enhance teacher awareness of small group work in the classroom. *IEEE Transactions on Learning Technologies*, 8(2), 187-200.
39. Matsuda, N., Cohen, W. W., & Koedinger, K. R. (2015). Teaching the teacher: tutoring SimStudent leads to more effective cognitive tutor authoring. *International Journal of Artificial Intelligence in Education*, 25(1), 1-34.

40. Mavrikis, M., Gutierrez-Santos, S., & Poulouvasilis, A. (2016). Design and evaluation of teacher assistance tools for exploratory learning environments. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge* (pp. 168-172).
41. Miller, W. L., Baker, R. S., Labrum, M. J., Petsche, K., Liu, Y. H., & Wagner, A. Z. (2015). Automated detection of proactive remediation by teachers in Reasoning Mind classrooms. In *Proceedings of the Fifth International Conference on Learning Analytics and Knowledge* (pp. 290-294).
42. Molenaar, I., Horvers, A., & Baker, R. S. (2019). Towards hybrid human-system regulation: Understanding childrens' SRL support needs in blended classrooms. In *Proceedings of the 9th International Conference on Learning Analytics & Knowledge* (pp. 471-480).
43. Newell, A. (1994). *Unified theories of cognition*. Harvard University Press.
44. Nilsson, N. J., & Nilsson, N. J. (1998). *Artificial intelligence: A new synthesis*. Morgan Kaufmann.
45. Nye, B. D. (2014). Barriers to ITS adoption: A systematic mapping study. In *International Conference on Intelligent Tutoring Systems* (pp. 583-590). Springer, Cham.
46. Ogan, A., Yarzebinski, E., Fernández, P., & Casas, I. (2015). Cognitive tutor use in Chile: Understanding classroom and lab culture. In *International Conference on Artificial Intelligence in Education* (pp. 318-327). Springer, Cham.
47. Olsen, J. (2017). Orchestrating Combined Collaborative and Individual Learning in the Classroom. *Unpublished doctoral dissertation, Carnegie Mellon University*.
48. Olsen, J. K., Aleven, V., & Rummel, N. (2015). Predicting Student Performance in a Collaborative Learning Environment. *International Educational Data Mining Society*.
49. Olsen, J., Rummel, N., & Aleven, V. (2018). Co-designing orchestration support for social plane transitions with teachers: Balancing automation and teacher autonomy. *International Society of the Learning Sciences, Inc. [ISLS]*.
50. Plass, J. (2020). Adaptive Learning | Gedanken Spiele. Available at: <http://janplass.com/index.php/2016/07/07/adaptive-learning/> [Accessed 20 Feb. 2020].
51. Prieto, L. P., Dlab, M. H., Gutiérrez, I., Abdulwahed, M., & Balid, W. (2011). Orchestrating technology enhanced learning: A literature review and a conceptual framework. *International Journal of Technology Enhanced Learning*, 3(6), 583.
52. Prieto, L. P., Sharma, K., & Dillenbourg, P. (2015). Studying teacher orchestration load in technology-enhanced classrooms. In *Design for Teaching and Learning in a Networked World* (pp. 268-281). Springer, Cham.
53. Ritter, S., Yudelson, M., Fancsali, S. E., & Berman, S. R. (2016). How mastery learning works at scale. In *Proceedings of the Third (2016) ACM Conference on Learning @ Scale* (pp. 71-79).
54. Rodríguez-Triana, M. J., Prieto, L. P., Martínez-Monés, A., Asensio-Pérez, J. I., & Dimitriadis, Y. (2018). The teacher in the loop: customizing multimodal learning analytics for blended learning. In *Proceedings of the 8th International Conference on Learning Analytics and Knowledge* (pp. 417-426).
55. Roll, I., Wiese, E. S., Long, Y., Aleven, V., & Koedinger, K. R. (2014). Tutoring self-and co-regulation with intelligent tutoring systems to help students acquire better learning skills. *Design recommendations for intelligent tutoring systems*, 2, 169-182.
56. Rummel, N. (2018). One framework to rule them all? Carrying forward the conversation started by Wise and Schwarz. *International Journal of Computer-Supported Collaborative Learning*, 13(1), 123-129.
57. Schoenfeld, A. H. (2010). *How we think: A theory of goal-oriented decision making and its educational applications*. Routledge.
58. Schofield, J. W., Eurich-Fulcer, R., & Britt, C. L. (1994). Teachers, computer tutors, and teaching: The artificially intelligent tutor as an agent for classroom change. *American Educational Research Journal*, 31(3), 579-607.

59. Sherin, M., Jacobs, V., & Philipp, R. (Eds.). (2011). *Mathematics teacher noticing: Seeing through teachers' eyes*. Routledge.
60. Simard, P. Y., Amershi, S., Chickering, D. M., Pelton, A. E., Ghorashi, S., Meek, C., Ramos, G., Suh, J., Verwey, J., Wang, M., & Wernsing, J. (2017). Machine teaching: A new paradigm for building machine learning systems. *arXiv preprint arXiv:1707.06742*.
61. Soller, A., Martínez, A., Jermann, P., & Muehlenbrock, M. (2005). From mirroring to guiding: A review of state of the art technology for supporting collaborative learning. *International Journal of Artificial Intelligence in Education*, 15(4), 261-290.
62. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
63. Tissenbaum, M., & Slotta, J. (2019). Supporting classroom orchestration with real-time feedback: A role for teacher dashboards and real-time agents. *International Journal of Computer-Supported Collaborative Learning*, 14(3), 325-351.
64. VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197-221.
65. VanLehn, K. (2016). Regulative loops, step loops and task loops. *International Journal of Artificial Intelligence in Education*, 26(1), 107-112.
66. VanLehn, K., Burkhardt, H., Cheema, S., Kang, S., Pead, D., Schoenfeld, A., & Wetzel, J. (2019). Can an orchestration system increase collaborative, productive struggle in teaching-by-eliciting classrooms?. *Interactive Learning Environments*, 1-19.
67. VanLehn, K., Cheema, S., Wetzel, J., & Pead, D. (2016). Some less obvious features of classroom orchestration systems. *Educational Technologies: Challenges, Applications, and Learning Outcomes*, 73-94.
68. van Leeuwen, A., & Rummel, N. (2019). Orchestration tools to support the teacher during student collaboration: a review. *Unterrichtswissenschaft*, 47(2), 143-158.
69. van Leeuwen, A., Rummel, N., & Van Gog, T. (2019). What information should CSCL teacher dashboards provide to help teachers interpret CSCL situations? *International Journal of Computer-Supported Collaborative Learning*, online preprint.
70. Walker, E., Rummel, N., & Koedinger, K. R. (2014). Adaptive intelligent support to improve peer tutoring in algebra. *International Journal of Artificial Intelligence in Education*, 24(1), 33-61.
71. Wickens, C. D., Gordon, S., Liu, Y., & Lee, J. (1998). *An introduction to human factors engineering*. New York: Longman.
72. Williams, J. J., Kim, J., Rafferty, A., Maldonado, S., Gajos, K. Z., Lasecki, W. S., & Heffernan, N. (2016). Axis: Generating explanations at scale with learnersourcing and machine learning. In *Proceedings of the Third (2016) ACM Conference on Learning@ Scale* (pp. 379-388).
73. Yacef, K. (2002). Intelligent teaching assistant systems. In *Proceedings of the 2002 International Conference on Computers in Education*. (pp. 136-140). IEEE.
74. Zhu, X., Singla, A., Zilles, S., & Rafferty, A. N. (2018). An overview of machine teaching. *arXiv preprint arXiv:1801.05927*.