

LEARNING TO TEACH: AN ARCHITECTURE FOR SELF-AWARE EDUCATIONAL GAMES

ABSTRACT

In this proposal I describe a new kind of educational game architecture: the self-aware educational game. Instead of simply amplifying an educational game designer's best guess at a teaching strategy, self-aware games will actively reason over uncertainty about how a subject should be taught, perform experiments, and meaningfully adapt to collected evidence on a global scale. Combining accepted educational theory, evidence-based educational practices, probabilistic inference carried out by machines, and deeply generative game design, self-aware games do not presume the effectiveness of any single teaching strategy; they learn to teach.

INTRODUCTION

Current educational games employ fixed models of pedagogical structure, missing opportunities to improve themselves in response to experience teaching.

In traditional educational games (games with static level progression), the game will continually present a learner who is stuck with the same challenge over and over. Removing such roadblocks with careful level design is very difficult and in many cases impossible (because what works for one learner may not work for another). Further, the information needed to make informed changes to level design is often not available until the game is finally deployed and a history of player data accumulates.

In the current trend toward adaptive educational games, a game system will attempt to tailor itself to the learner's preference and learning style. These games realize measurable benefits for their players over traditional educational games, but they are blind to global roadblocks in which whole populations of learners get stuck or fail to benefit from the system's adaptively chosen exercise or intervention. Player-adapting educational games respect the uncertainty we have about a learner's progress, but fail to capture uncertainty in how a curriculum should be modeled and how the system's available interventions relate to the curriculum on a global scale.

Where traditional educational games are designed with a one-size-fits all pedagogy, adaptive games replace this with an equally self-confident model that assumes all variability in the best teaching strategy is accounted for by individual differences between the learners. In educational theory, the use of evidence-based practice (EBP) attempts to replace the use of intuition, tradition, and rules of thumb with strategies that are validated by recorded evidence and are open to change as new evidence is made available. EBP is stymied in classroom adoption by confounding effects (such as variable teacher skill, availability of funding, and varying student backgrounds), limiting the transferability of evidence between scenarios. Lacking a large body of evidence in any particular domain, the use of universal and globally-inflexible teaching strategies continues. Thus, traditional and adaptive educational games are a natural outgrowth of the current educational climate.

Computer-mediated games have qualitatively different capabilities than the traditional classroom and human teacher, and thus are not subject to the same constraints that halt adoption of EBP. Because a computational system can be made to interpret a teaching strategy in the exact same way across any number of students and record its observations in a common format, the potential for evidence gathering, sharing, and exploitation in computer-mediated educational games far outstrips the ability for individual teachers in loose association to do the same. A machine's strict adherence to an objectively defined strategy and documentation policy makes it an ideal practitioner of EBP, removing many (but not all) of the confounding effects which block EBP adoption.

I recognize that one of the other major factors blocking EBP is the overly mechanistic nature of teaching it suggests. In a world where educational computer games took care of low-level compliance with EBP strategies, human teachers would be free to work more creatively, finding and exploring gaps in the curriculum, and inventing new strategies that could be directly compared against and eventually integrated into the evidence-based practice. Games employing EBP will compliment the human teacher's natural strengths, empowering teachers as innovators instead of cogs in an educational machine made of people.

In my vision for self-aware educational games, I propose the development of systems that reason over the presence (and lack) of evidence for the effectiveness of various exercises and interventions. These systems will realize educational gameplay that is adapted both to differences between individual learners and differences between the assumptions used to initialize the system's overall pedagogical strategy and evidence collected across the entire player population. When a self-aware educational game is failing to produce the desired impact on learning, it will know it, and it will be able to report precise statistics on what it actually was able to achieve. Self-aware games will actively participate in curricular design, producing and consuming evidence for the practice of alternative strategies (within a pre-set scheme of variations), instead of simply being vehicles to convey a designer's a-priori beliefs about how a particular topic should be taught.

EDUCATIONAL BACKGROUND

Before launching into my proposal for an architecture for self-aware educational games, I want to review a few key ideas from educational theory to motivate some of my architectural choices.

ZPD

Lev Vygotsky, a noted psychologist and early educational theorist, defined the idea of the "zone of proximal development" (ZPD), in which he suggested all of a learner's experiences should fall for effective development. The ZPD is "the distance between the actual developmental level as determined by independent problem solving and the level of potential development as determined through problem solving under adult guidance, or in collaboration with more capable peers". While the ZPD was originally conceived in the context of child development, it is applicable to learners of all ages as the idea of tracking both the skill a learner can demonstrate alone and the fringe of skills they can demonstrate only with external assistance.

Current adaptive educational games track only one boundary of the ZPD: those skills a learner can demonstrate alone. When collecting and integrating evidence on the effectiveness of various interventions, it is not enough to know whether a student had the apparent prerequisites for the intervention in question (with respect to a tentative curriculum model), knowing whether the intervention falls within the learner's ZPD is critical.

In game design, Koster's "theory of fun" and (from psychology) Csikszentmihalyi's theory of "flow" are both concepts that are deeply related (but not equivalent to) Vygotsky's ZPD. All of these ideas revolve around giving a

player/learner experiences at the fringe of their abilities. While flow is difficult to measure (except via very indirect proxy measures), the ZPD can be directly measured through the solo and collaborative exercises already present in an educational game's play.

SCAFFOLDING

The concept of scaffolding (which actually emerged from the application of the ZPD to education) is the use of assistance by a teacher or more confident peer (and eventually a computational system) to give a learner experiences inside of their ZPD. Though assisted experiences, a learner can take on problems which she could not tackle alone and eventually accumulate enough understanding to work these unassisted. The practice of scaffolding allows a teacher to synthesize exercises with fine-tuned levels of challenge even if the target domain does not have a naturally smooth progression.

Uncertainty about how and when to apply scaffolding to best accelerate a learner's progress represents is an area where self-aware games should place the bulk of their introspective focus. The production and application of evidence-based practices as a result of the large-scale experiments automatically carried out by these self-aware games is the primary element distinguishing them from past educational games. A scaffolding strategy which is sensitive to both the global statistics of the player population and individual differences should provide massive benefit for both the learners and external reviewers who will look in on the evidence collected.

RECIPROCAL TEACHING

One specific form of scaffolding is the use of a learning peer to play the role of the teacher; this is reciprocal teaching. Reciprocal teaching not only scaffolds the experience for the player taking the role of the learner, but it gives the player taking the role of the teacher an experience that is often absent from both the classroom setting and educational games alike. In teaching (or at least moderately assisting) another, the teacher comes to understand the material at a much deeper level, building both competence and confidence in the subject.

Self-aware games will know that the palette of available exercises it has to offer learners will not provide a flawless method for educating students (they will be continuously updating the statistics that confirm this). Part of what these games should model is the benefit provided by out-of-game scaffolding provided by the interaction with peers. By recording when learners are placed in situations must help each other and accounting for knowledge transfer outside of in-game actions, a self-aware system's pedagogical model will enable reflection on the relative effectiveness of its exercises with respect to this out-of-band learning. Thus, self-aware systems will also know when it is best to stop hand-holding a stuck learner with small exercises and when to give a group of learners a larger challenge that can specifically engage them in reciprocal teaching to engage the stuck learner in a different modality.

ARCHITECTURE

Thus far, I have suggested that self-aware game systems will have some type pedagogical model that includes uncertainty, a palette of exercises and interventions, a means of synthesizing the appropriate exercise for the learner at any time, and a means to monitor the students progress (really, their ZPD). In this section, I will flesh out these elements of the system architecture with more detailed examples.

PEDAGOGICAL MODELS

I imagine self-aware educational games primarily maintaining three types of probabilistic models over pedagogical uncertainty. A player knowledge model will continually track individual players' ZPD details, an intervention model will track the effectiveness of various interventions assuming the requirements for applying that intervention are met, and a curriculum model will track the apparent dependence of topic elements on one another and the dependence of interventions on those topic elements.

CURRICULUM MODEL

Working backwards, the curriculum model represents a system's understanding of the relation between topics and teaching strategies in a subject. In a high-school mathematics context, the curriculum model might contain the (probabilistically hedged) assumption that knowledge of prime factorization depends on knowledge of advanced fractions and exponent notation. It would also describe the assumption that a particular scheme of exercise, perhaps determining the list of prime unique factors for an integer less than 1000 without regard for cardinality, hinges on a basic level of understanding of prime factorization and will produce an intermediate-grade understanding of the topic if completed successfully.

The curriculum model encodes a picture of how a subject is structured and what the available methods are for teaching it (along with the conditions for their availability in terms of an individual's progress). It does this without reference to any particular learners or any particular experiments applying interventions. Evidence coming from actual experiments having particular students attempt particular exercises in particular areas of their ZPD will play into improving the accuracy and certainty of the curriculum model, but only indirectly through relatively advanced statistical inference (which is generally beyond the realm of un-assisted human calculation). Machine learning and inference tools will perform maintenance on the curriculum model. Intuitions and traditions for a subject can be seamlessly integrated into this evidence-based framework by providing initial/prior distributions and weights for the model's structure and parameters.

INTERVENTION MODEL

Moving closer to the evidence that a self-aware game system will actually record, the intervention model is a description of the observed effectiveness of the system's available exercises and interventions. Because an intervention can fail in many ways (it can simply be a poor intervention, it could be misplaced in the curriculum, or it could be a poor match for a student's learning state), it is important to have a very clear semantics for the intervention model. In the architecture I am proposing, the intervention model should reflect effectiveness (chance of producing an intervention's stated goal) of a method assuming that method's requirements (as described in the curriculum model) are fully satisfied.

The intervention model is intended describe a kind of ideal situation which is never actually cleanly realized in practice. Practically, a teacher is never quite sure of a learner's knowledge, nor is he sure of the method's exact requirements. Concretely, this means observing that a single application of an intervention fails does not completely count against that method; it means that to the degree that the models were certain that the student was ready and that the method was correctly applied, the method's modeled effectiveness should be fractionally reduced. Making the intervention model a readable and trustable resource requires reasoning over the joint uncertainty in the system's other models, again, a task well beyond casual mental calculation.

PLAYER KNOWLEDGE MODEL

Finally, the player knowledge model most closely resembles something used in current adaptive educational games. Where my proposal differs from a simple model accounting for what the player knows is in its use of the ZPD and its unique method of update. Instead of recording whether a learner does or does not understand a particular micro-topic, the player knowledge model will (again, with probabilistic hedging) describe a learner's ability to demonstrate a particular skill with different levels of scaffolding (assistance). In this way, the player knowledge model actually describes a detailed map of the learner's ZPD.

The player knowledge model is updated, not by simply checking off topics as the learner demonstrates a skill once, but in a global probabilistic inference pass that accounts for the structure and uncertainty in the other two models. As an example, if a self-aware game is highly confident in both its curriculum and intervention models (unlikely, but possible), upon giving a particular learner an exercise that they observably fail, the system may decide to believe that the student actually did learn the desired skill, but it was a fluke in evaluation that best accounts for the observed failure in that single instance.

In comparison to a game's global population of players, relatively little evidence is actually scoped to any particular learner. Global probabilistic inference provides a sound means to transfer observations about one student to another. In a way, this permits the self-aware game system to (partially) believe the over-general idea "if you give the exercise, the students will learn the topic" while having hard data to support the exact degree to which this general assumption is applied and balanced with the sparse observations for a particular learner.

EXERCISES AND INTERVENTIONS

Self-aware educational games will, to a degree, function as both automated game designers and active game masters (in the table-top gaming sense). Instead of presenting a static sequence of preset levels, a self-aware game will know how to best assemble an exercise for the current players and when to intervene to steer gameplay in a new direction.

The primary technology that will enable this mode of gameplay is the reconfigurable exercise. Building on game design research in procedural content generation, self-aware games will assemble recombinable parts of game objects, levels, and stories, adjusting free parameters and structures in accordance with the pedagogical models.

Building on the kinds of exercises and interventions already used for a particular subject (i.e. those used in textbooks and teaching manuals), reconfigurable versions of these methods should be created with a rich space of variation for how much scaffolding is embedded (how much assistance is provided as part of the method's setup and execution). Returning to the high-school math context, a simple problem involving finding solutions to a quadratic equation can be scaled in its difficulty by changing the level of specificity of input (free-form input, multiple-choice input, picking the name of a solution strategy from a list, or simply deciding if the problem is one that is solvable with a named technique), or the level of accuracy required (should the answer be numerically correct, be within a certain range, or simply start off with the right step).

Another important way in which exercises and interventions should be reconfigurable is in how they exploit opportunities for reciprocal teaching. A larger problem can be broken into steps which will be solved by different learners working together. A good factoring of problems into subproblems should permit combinations of subproblems which are dependent and asymmetric, allowing the self-aware system to use the same overall problem to work with very different individual learners working in concert and motivating them to work out some of the concepts verbally, outside of the game's core mechanics. For a quadratic equation solving problem, one learner would be asked to factor the equations and the other to solve the equation given the factoring from the

first, exploiting the feeling of ownership over intermediate results and the reinforcing benefits of successful cooperation. Where a joint exercise is failed, the system will understand it to have formally failed for both learners (with the effects of learning from failure already being accounted for in the intervention model).

Determining the set of reconfigurable exercises and interventions made available to a self-aware educational game is itself a difficult design problem. The burden of having to produce a set that will be a-priori effective is, through this architecture, reduced because the system will work to record and report the effectiveness of each method. Giving a self-aware system a palette of exercises that is overly expressive will result in the system gathering little evidence for each particular type of exercise (because, acknowledging its own uncertainty, it will eventually try all alternatives). Thus, the best palette is one which gives the system room to explore while not delaying too long the time to high-confidence feedback on dependence and effectiveness.

SYNTHESIS OF GAMEPLAY ELEMENTS

Given the system's pedagogical models and a palette of reconfigurable exercises and interventions, there is always an optimal action that the system should take next to optimize some metric over the player population (such as the sum total of all learner's skills or perhaps the progress of the weakest learner -- "no child left behind"). While this optimal action is mathematically defined, it is both unlikely that this optimum could be found algorithmically in reasonable time (due to complexities in the interaction between the pedagogical models and the reconfigurable exercise palette) and relatively unimportant to actually produce. Given honest (probabilistically correct) modeling, it seems unlikely that a self-aware would ever have very strong confidence in any single next exercise (given the very noisy nature of education when viewed mathematically). What is important is only that the system chooses a next exercise in which it has a reasonable level of confidence relative to the alternatives.

Methods selected for experiment early in the deployment a self-aware game, even if they were selected on the basis of a very poorly tuned pedagogical model, still contribute meaningful evidence which helps build the system's confidence in alternate tunings of the model. In the architecture I imagine, models will be updated using all available experimental data, not just the most recent data. This corresponds to the assumption that some subject, perhaps high-school mathematics, does not fundamentally change during the deployment of the game.

Beyond simply reasoning over the best next-step, self-aware educational games can introspect on their pedagogical models to make longer term plans for exercises. A system may decide (and declare to the learners up front, to give them context) "we are going to three easy exercises of type-X and then one group exercise of type-Y at the end to make sure you've got the idea". Where curriculum-space plans may seem dry from the perspective of a direct reading by the learner, these same plans could be communicated to a secondary system which is generating coherent narrative elements for the game. Accurate likelihood annotations on such plans could advise a narrative generation system to perform contingency planning for a particular step in the intended story that the teaching system thinks the players are most likely to fail at.

Using the continually-update pedagogical model as a means to forward-simulate player learning in a statistically honest way, it should be possible to solve for coherent sets of exercises that take into account player skill, play history, external thematic elements, and social relations between players that robustly achieve learning goals. While it is tempting to consider generating whole missions up-front in this way, it would be both technically challenging and serve to defeat one of the strengths of self-aware systems: their ability to transform in response to local and global evidence gathered from play experience. Self-aware educational games will be at their best when they are agile, given the ability to rapidly reconfigure the current game setup.

MONITORING PROGRESS

Dynamic assessment (another Vygotskian concept) is the idea of interleaving assessment of a learner's state with the interventions used to teach them. Self-aware games, with a dramatically richer embedded pedagogical model available for introspection than traditional games, should be able to improve their understanding of a learner's ZPD-state from almost any computer-mediated interaction the learner has in the game, not just from isolated pre- and post-tests that occur outside of the game's core context. Imagine that a learner who had been failing to demonstrate a certain skill has just successfully completed a joint exercise in which a player in their role usually possesses the skill in question. We can't know for certain whether the learner just acquired the skill during this last exercise or if it was their partner that covered for them, but the estimated probability that they did indeed freshly acquire the skill will fractionally contribute to that learner's player knowledge model. This indirect inference does not need to be specially engineered into a monitoring system; instead it falls out of using the same kind of inference tools over a large model constructed in a common probabilistic logic framework.

IMPLEMENTATION SUGGESTIONS

There are, of course, many ways to instantiate this architecture for self-aware educational games using different representations and algorithms. In this section, I will share my initial thinking on how to realize the required probabilistic inference and synthesis of game elements.

PROBABILISTIC INFERENCE

A promising avenue for representing probabilistic models over the kind of pedagogical structures I described above is the framework of Markov logic networks (MLNs). Markov logic generalizes first-order logic with probabilistic inference. Where a traditional first-order logic model will state rules and facts with absolute certainty, MLNs possess weights encoding the confidence of various rules (allowing consideration of competing, conflicting hypotheses). The weights in a MLN can be used to compute probabilities for logical facts. The mapping from weights to probabilities is non-trivial and sometimes difficult to interpret, but this does not come into play with the proposed game architecture.

MLN-based *Alchemy* is an actively maintained, open source tool that implements statistical relational learning and probabilistic inference. Specifically, it implements all of procedures needed to update and apply the proposed pedagogical models. *Alchemy* has already been used to learn models and make predictions for social network analysis, web-link structures, and natural language sense disambiguation. All of these involve the same kind of uncertainty over large graph structures that I expect to arise in the implementation of the pedagogical models.

MLNs are not a complete answer for implementing the pedagogical models, of course. MLNs provide only a substrate for modeling logical relations of variable certainty. To my knowledge, the dependence structure for a subject like high-school mathematics has never been written in a formal representation that admitted any form of uncertainty. Simply constructing an initial model in this representation (perhaps using published curricula as a starting point) and hand-modifying weights to relax certainty would represent an original research result in and of itself. Towards making a self-aware educational game, the network used to initialize the system need not be at all precisely tuned, it just needs to give non-zero confidence to a more accurate configuration of weights and structure so that the system may eventually discover it.

Of the models inside of the overall pedagogical model, the curriculum model seems likely to be both the most complex and the one with the most evidence which must be aggregated. Without experimentation, I can't predict

in objective terms how big the MLN for the curriculum model will be, but I can imagine that it might be too big to be updated regularly. To create a self-aware game system that is sensitive on both a global and an individual scale, it may be necessary to hold the curriculum and intervention models constant while performing only local updates to individual player knowledge models as local gameplay progresses. On a larger time scale, with access to statistics reported from remote deployments, and with access to bulk computation resources, the curriculum and intervention models can be updated centrally, with the results infrequently broadcast to deployed game systems. This strategy encodes the idea that player knowledge models will be in continuous flux while curriculum and intervention models are likely to stay on a slow trajectory of convergence to a fixed, global truth. The pedagogical model will be expected to shift on a time scale that matches the rate at which new exercises and interventions are made available to the system (materially changing the pedagogical landscape of the subject).

SYNTHESIS METHODS

At any time, the combined pedagogical model will induce a probabilistic distribution over what exercises and interventions will cause a learner to progress. For a relatively simple space of reconfigurable exercises (one that could also be expressed in Alchemy's input language), it should be possible to use the existing reasoning tools to directly sample exercises from this distribution. This approach elegantly solves the problem of exploration of alternatives vs. exploitation of the best strategy -- the approach simply samples exercises with a probability proportional to their chance of having the desired effect for the learner in question.

For schemes of exercises and interventions that are impractical to model in logic, it should be possible to sample exercise designs from an external generator and use Alchemy only to score them (with a learner-progression-success probability). In this way the pedagogical models could be used to implement the black-box fitness function which a metaheuristic search method would use to select a high-fitness exercise. When consider this option, it is important to keep in mind that the fitness function will be changing as a result of evidence gathered each time the learner completes an exercise (their ZPD might have significantly shifted during the most recent exercise). Any synthesis method chosen should be able to produce a reasonable next-step exercise within the time between the completion of one exercise and the start of another if it is going to allow the enclosing self-aware game maximum agility with respect to learner progress.

CONCLUSION

I have proposed an architecture for self-aware educational games, games that reason over their own uncertainty in pedagogical matters and react to both individual and global patterns in an evidence-based manner. This architecture is qualitatively different than that of the adaptive and traditional (static progression) educational games that came before it. It makes pervasive use of probabilistic inference by machines to implement the rigorous scientific demanded of but resisted by educators in the traditional classroom setting. It is my belief that self-aware games educational games will bring clarity, objectivity, and transparency to curriculum and intervention design while celebrating the creativity of the human innovators who produce the new teaching methods which these self-aware systems will select for experimentation.

Adam M. Smith (amsmith@soe.ucsc.edu)

April 23, 2011