The Daemon as Educator: Ubiquitous Access to A Personal Mentor

Bernard Rosell
AT&T Laboratories,
200 Laurel Avenue, Middletown, NJ 07748 USA

and

Raluca Sirbu
AT&T Laboratories,
200 Laurel Avenue, Middletown, NJ 07748 USA

ABSTRACT

In Socratic philosophy, human beings had access to an inner voice that could be counted on to provide guidance and assistance when needed. This inner voice or daemon was always available and had access to realms of knowledge that mortal consciousness did not have. We propose that the field of Machine Learning has reached the level of maturity where the construction of a personal daemon is now feasible. The ubiquity of broadband access to the internet can guarantee that this personal daemon is always available. Such a construct can greatly increase the efficacy and availability of education for all citizens.

Keywords: Cloud Computing, Machine Learning, User Modeling, Personalized Mentor, Educational Adjunct

1. INTRODUCTION

Classical Greek philosophers developed the idea of daemons (δαημονεν), divine entities that attach to individual human beings, with the purpose of guiding them to grow and mature into their fundamental nature and character. Towards the end of the Republic [1], we read "All the souls had now chosen their lives, and they went in the order of their choices to Lachesis, who sent with them the *genius* whom they had severally chosen, to be the *guardian of their lives and the fulfiller of the choice:.."* (Italics ours). Each daemon accompanies its charge throughout their entire life and constantly directs its charge to fulfill their potential. The concept of the daemon as one's personal companion and guide found its greatest expression in the life of Socrates, who credited his daemon as the source of his wisdom and courage in the face of death.

Suspend disbelief for a moment as we recast this wonderful idea in the contemporary idiom. The phrase "choice of life" we can understand as the genetic and social matrix of an individual. But what can we understand by the idea of the daemon? Implicit in the idea are the following key constructs.

1. The daemon possesses a deep understanding of the human to which it is attached
2. The daemon has access to a realm of knowledge no individual human being could attain.
3. The daemon is an advocate for its charge, i.e, it can only act in ways that positively affect the development of its charge
4. The daemon is always available

We propose that a confluence of developments has now made it possible to create an entity that embodies all these constructs and to enlist that entity in the service of humanity compatible with the vision of those ancient Greeks philosophers. The developments we refer to are respectively, the introduction of cloud computing, high bandwidth access to the Internet from mobile devices and significant developments in the field of machine learning. Specifically, we propose the development of a software agent running on a virtual machine implemented in the Internet (clouds) that has access to all the data on the Internet and access to the most sophisticated machine learning tools. This process will attach to individual human beings, learn all about them and then guide them in developing their full potential. In this paper, we focus our attention on the ramification of this idea to education. In the following sections, we argue that such a program is indeed feasible and that the time is ripe to implement it.

All 4 elements listed above are essential, though from our perspective items 1 and 3 are the most controversial and require some justification. Most significant is the claim that an artificial entity, i.e., a computer process, can come to know and understand a human being and then significantly impact their lives in a positive manner. The lynchpin of the whole argument is the idea that a computer program can come to "know" a human being in the sense that we would say person A "knows" person B very well.

Consider the now quite common scenario of logging into an on-line bookseller or on-demand movie site and being presented with a list of recommendations. Chances are good that a reasonable number of the presented recommendations were for the most part interesting to you and were titles that you either had already bought or would consider purchasing. After taxing yourself deciding whether to go ahead and purchase any of these recommendations, perhaps you turn on Pandora to relax to some of your favorite music. On the surface, these applications really seem to have learned something meaningful about you and to be able to use that knowledge in a way that improves your life. Recent advances in Recommender Systems [14] rely on sophisticated algorithms that analyze the similarities between previously selected items. It is essential to note that to be effective the algorithm would have had to have previous interactions with you. As with all inductive algorithms, the more interactions the more likely the algorithm is to provide a valid recommendation. To determine the set of musical selections most likely to please you, Pandora analyzes the musical structure of the songs you like and then selects those songs with similar structure. The reference information is contained in the Music Genome, created by a set of musicologists and containing more that 10,000 pieces of music codified with respect to hundreds of musical attributes such as harmony, rhythm, tonality, range, etc. [12, 20]. The key item here is that the algorithm has access to a rich set of reference data.

When we say that a software process is an agent we require that it can learn and that it acts to achieve some goal. The Automated Travel Assistant [18] provides an excellent example of an application that can learn an individual's travel preference, search airline timetables and provide a set of tailored recommendations. Indeed, the study of intelligent agents has brought forth a myriad of applications that support or automate human functions. Neural networks were used to craft a personal news service based on an explicitly learned user model [8]. We argue that it is now feasible to extend this approach to create an intelligent advocate that will monitor and support the educational development of its charges. This Educational Daemon will take advantage of all the breakthroughs in machine learning, user modeling and developmental psychology to provide a supportive guide that is always accessible, easy to interact with, and tailored to each individual student. It is personal in the sense that it will actively learn the students it supports.

Second generation "intelligent agents" are very sophisticated machines. Due to advances in computer hardware, distributed processing and machine learning the computational power of these agents is significantly expanded. These machines can automate human tasks, and they can interact with people in human-like ways that include natural language and the ability to participate in a conversation or dialog [9, 5].

Building such a personal entity, **Y**our **E**ducational **D**aemon (YED) is clearly an ambitious project, one that will require the collaboration of experts on multiple fields. Clearly, a key competence of this agent would be the ability to understand or model the student it will serve. The following description from the mission statement of the Journal of User Modeling aptly summarizes the areas of research and development that will be required to bring this idea to reality

"……conceptual models, mental models; levels of user expertise; intelligent information retrieval; adaptive hypertext and hypermedia; adaptation to the handicapped and elderly; user stereotypes; formal representation of user and student models; applications in office machines and consumer electronics; and privacy and security of information for personalization.

User Modeling and User-Adapted Interaction is ideal for researchers, students and industrial practitioners in artificial intelligence, human-computer interaction, linguistics, and the instructional sciences"

2. DEFINING THE DAEMON

In the introduction we mentioned 4 key characteristics required of our daemon: (1) developing a deep understanding, i.e., learning about the student that it will serve; (2) determining the target environment in which the student will operate viia access to reference data; (3) delivering a set of training /learning regimens tailored to the individual student and to the target environment and (4) being always available. It should be noted that theses characteristics are not unrelated. Actions such as playing a game with a child can serve a tutoring function as well as provide an assessment of the child's specific cognitive and social skills. Such activities provide the daemon with useful information in addition to increasing the child's self confidence [29, 30].

In essence, YED is a learning advocate for the child it serves. Furthermore, this advocacy will evolve over time as the child moves though the education system. For the present we arbitrarily limit our focus to grades K – 12. This is clearly a formidable task but a feasible one as we hope to demonstrate.

It is essential that YED is always available to the student. For our purposes we consider the daemon as a background process running on a virtual machine. A new instance of this process is forked (created) when a new user (prospective student) subscribes to this service (attempts to communicate with their daemon) for the first time[1]. Since the daemon is always available and does not halt when the student detaches or closes their channel, it must preserve state information. In other words, the next time the student connects to it, the ED will start from where it last left off and incorporate all the learning from the last encounter.

We explore each of theses factors in more detail below.

2.1 Learning the Student

In the rich literature on machine learning [11], the breadth of concepts and/or skills that a machine can quantitatively be shown to learn run the gamut from abstract concepts such as Boolean expressions to playing games such as chess to driving a car. Here the target the YED must learn is the student: their strengths, weaknesses, preferences and interests. Initially the YED views the student as a tabula rasa and the information it obtains would be from the parent or teacher. As soon as the

[1] Processes can become daemons by forking a child process and then having their parent process immediately exit, thus causing init to adopt the child processes. The resulting process is not attached to a specific TTY (user) or Socket (interface) and does not die when the specific TTY or Socket session terminates.

student is capable of accessing their daemon, the YED will learn from these interactions. These interactions would take the form of:

1. Playing games with the student
2. Administering tests
3. Asking and answering questions or just having a conversation.

With each interaction the daemon builds a richer internal model of the student. In conjunction with externally provided data such as results of classroom administered tests or feedback from parental feedback, the internal model converges with increasing interactions to an accurate model of the student.

Of course we must be able to validate that the daemon can actually learn the child to which it is attached. While a child's educational profile is not as well defined as a class of Boolean expressions (e.g. k-term CNF formulas), the framework of the Probably Approximately Correct (PAC) learning model [17] provides a reasonable way to think about the problem.

In our construction, the daemon will have access to the full set of learning algorithms that have would have previously been shown to be effective and the issue of validation will have already been closed for each specific and measurable capability. For example, machine learning models have been demonstrated to effectively evaluate a non-native speaker's grasp of pronunciation and reading comprehension when learning a new language [36, 37 38].

However, it would seem reasonable to expect that an artificial entity should be held to the same standards of proof as I would if we both claimed to know someone. But what does it mean to "know someone"? My claim to know a specific person very well could be plausibly made if I could answer a sufficient number of verifiable statements that describe:

1. What that person is like i.e., their specific interests, abilities, personal style, goals and motivations [34].
2. How that person would act / perform in specific circumstance, e.g. "what would that person do if …" or "how would that person react if ….". For our purposes we are primarily concerned here with the level of skill acquisition [36].

Some examples of statements in the second category applied to a specific child named Bard would take the form:

- Bard will score in the 4^{th} quartile on the mathematics section of the TerraNova test
- Given a list of 10 words appropriate to Bard's grade in school, Bard will be able to spell at least 8 correctly?"
- Bard asks a lot of questions when exposed to a new ideas or situations?
- Bard reads at the 6^{th} grade level
- Bard wants to work with animals
- Bard suffers from asthma

We first assume that to learn a child in the context we are interested in is equivalent to learning the specific values of a potentially large but finite set (bounded in advance) of attributes. Such attributes would encapsulate personality traits, individual preference, values, important medical conditions, current educational attainment, and of course, cognitive strengths and areas for improvement. We call this set of attributes the child profile and note that it will evolve over time. This attribute set must satisfy three properties to be useful; it must be complete, predictive and computable.

1. The set of attributes must be rich enough to encapsulate a complete view of the child personality with respect to their educational development.
2. The set of attributes should be rich enough to allow a cadre of experts to predict how the child will respond in a specified situation or interaction.
3. The value of each attribute can be learned through interactions of the type we are considering here.

For each student **s** who is mentored by daemon $Y_{\mathbf{s}}$ let $I_{<s,t>}$ denote the set of all interactions that specific student of age t engaged in over some interval. Let $C_{<s,t>}$ be the set of all possible assertions that can be made about those interactions with respect to their outcomes and a label indicating the truth value of each statement. Examples would be scores on standard tests, reactions to structured social situations and first-person narrations as exemplified in the bullet list above. What matters is that each outcome is verifiable by an objective measure or via the direct observation of the child, parent, educator, etc.

We proceed as follows. Let D_S be an arbitrary distribution over the set of all students that have been mentored by a daemon and let $D_{C<s,t>}$ be an arbitrary distribution over $C_{<s,t>}$. Two parameters δ and ε with ($0 < \delta < 1$) and ($0 < \varepsilon < 1$), representing our bound on an acceptable error rate and required confidence level are chosen in advance. A set of students is chosen randomly according to D_S from each age group and for each student s we run the following experiment.

- Choose a list of statements $L_{<s,t>}$ drawn at random from $C_{<s,t>}$ according to $D_{C<s,t>}$. The scenario is provided to the daemon without the label
- The daemon $Y_{\mathbf{s}}$ will provide its prediction of how student **s** would respond / perform /act in each situation, i.e. will evaluate the truth of the assertion based on its knowledge of **s**.
- *err* $Y_{\mathbf{s}}$, the error rate of daemon $Y_{\mathbf{s}}$ for this student is calculated. $Y_{\mathbf{s}}$ makes an error if its prediction does not agree with the observed results, i.e., the truth values assigned $Y_{\mathbf{s}}$ by does not agree with the label of the assertion.

Then, we can reasonably say that a daemon can "learn a student" if

$$\forall \text{ t in the Index set}$$

EQ (1) $$P_{D_S}(err\ Y_{\mathbf{s}} \text{ on } L_{<s,t>} < \varepsilon) > 1 - \delta$$

In essence, for a student selected a random we can expect with probability at least 1- ε that our daemon will be able to answer at least 100 * (1- ε) % of the questions correctly.

Alternatively, we might require a stronger condition and require that for every child at every time period the algorithm must have a minimum level of competence. We would then conduct an experiment that was sufficiently powered to test if the error rate of the algorithm is < ε. In other words,

$$\forall \text{ t in the Index set } \& \ \forall \text{ s in S}$$

$$\textit{err } Y_{\mathbf{s}} \text{ on } L_{<s,t>} < \varepsilon \qquad \text{EQ (2)}$$

This constraint guarantees that for every child at every time period the algorithm must have a minimum level of competence. The real point of this discussion is that given a well defined set of attributes meeting the three criteria defined above the daemon could be constructed, component by component to satisfy any reasonable definition of competency.

Note that we are only addressing feasibility here and not performance. The issue of efficient learning is critical but outside the scope of this paper. However, the issue of efficient learning is strong related to 2 factors, i.e., the size of the profile and the amount of interactions with the child.

At first glance it might seem that the size of the child profile is daunting, but there are good reasons to believe that a small set of attributes or traits has strong predictive power. In one longitudinal study [32], 4 childhood traits were shown to predict adult behavior 40 years later. These traits, namely, verbal fluency, adaptability, degree of impulsivity and self minimization, can be learned by interaction with the child. We can generalize the verbal fluency trait to include texting and as we will note in section 2.4, the daemon will have ample opportunity to "speak" or text with its mentee and to observe its interactions with others. Other research has focused on the "Big Five" personality traits [33] and modifications that include self-directed goals and motivations [34]. The ubiquity of access to the cloud and hence the ease of access to the daemon and the wealth of potential interactions provide an ideal vehicle for the daemon to evaluate the child's interests and personal style. In the realm of machine learning, it has been shown that for a many of the problems studied a small o number of attributes are all that is necessary to make reasonable classifiers [39].

While there are no empirical results on complex agents of the type we are discussing, there have been significant results on the use of machine learning algorithms to implements many of the functions YED would require. Three areas stand out:

1. Test construction and analysis
2. Playing games
3. Concept learning

2.1.1 Test construction: A valuable test will allow the examiner to accurately categorize the ability of the examinee based on the overall performance of each item in the test. This is crucial in accurately determining what remedial steps are required and which enhancements could be offered. The same toolset will allow the daemon to assess the level of progress after a specific course of action has been completed. In addition, an effective test will help the student clearly see their strengths and areas for improvement [31].

The theory of test construction is concerned with determining the optimal type and number of items, (tasks or questions) will simultaneously maximizing the accuracy of the test Stated this way, the problem becomes an optimization problem and hence a candidate for mathematical based methods [7, 10, 2, 18] and machine learning methods [15, 6]. Adaptive techniques enable the creation of dynamic assessments in which the question generation function selects subsequent questions based on the difficulty of the previous questions and the performance of the student on those questions [27]. This approach can increase the efficiency of the testing process as well as provide the student with a quick self assessment. Once YED has selected the appropriate test(s) for a given set of developmental targets and evaluated the results, the internal times series model represent the student is updated to reflect the new information.

2.1.2 Playing Games: There is a rich body of research in machine learning on the ability of algorithms to learn to play games and to learn from games [19, 24, 25, 30]. Our interest is in this modality is twofold: first to evaluate the student within the developmental life-cycle and second, to use games as a tool to develop perceptual skills, motor skills and cooperative behavior. We can imagine the daemon playing a game of 20 questions or GHOST with the student and using the interaction to evaluate the student's ability for logical thinking and vocabulary aptitude.

Videogames are a particularly fruitful tool to engage children (and adults for that matter).There already exist game development methodologies that allow the results of the game to be made available to a Learning Management System (LMS) Specific targets can be encoded in the game and then made available in assessment logs or reports to the LMS [25], or in our case, to the daemon.

2.1.3 Concept Learning: There is a robust body of theoretical research in the field of machine learning on concept learning [3, 4, 13] A large variety of concepts such as regular sets, decision lists and many classes of Boolean expressions have been shown to be efficiently learnable if the learning algorithm is able to query an oracle (teacher). The query can be either a proposed hypothesis or an example. In the first case the teacher responds in the affirmative or provides a counterexample. In the second, the oracle responds yes if the concept would hold true for the example, and no otherwise.

In some sense, of course, this is how children learn concepts as well. An algorithm would be able to judge a student's grasp of a specific concept by means of similar interactions, i.e., asking the students for examples and counterexamples or by providing them to the student and asking if they are valid.

A key heuristic in concept learning is to reduce complex concepts into simpler components, learn the simpler components, and then aggregate these back into the initial complex concept. [27]. YED would use a similar approach. A student interacting with YED would be presented examples that implicitly build a "concept hierarchy" or inheritance structure to represent the complex concept.

The key idea here is that we can exploit the growing knowledge of how machines can efficiently learn concepts to building machines that can efficiently evaluate whether a human understands the same concepts. As with user modeling and intelligent agents in general, an inter-disciplinary approach is needed to define the optimal set of human-daemon interactions.

We would be remiss if we did not mention that machines can learn to recognize emotional constructs and states as well [28]. This ability will be a core competence of the daemon in its interactions with its charge.

2.2 Access to Reference Data

Analogous to the Socratic deamon, YED is able to access information extending beyond the individual it serves. In order to effectively develop an effective training plan, the daemon must be able to relate the student's capabilities to the appropriate development target. This target is a function of the student's chronological age and developmental profile. The curriculum requirements are a key source for this information. Since these requirements are specific to geography, and highly dynamic, they would be best accessed from the appropriate repository. As new games or tests became available, the daemon would incorporate them into its interactions with the student.

In addition, the daemon would have access to a wide variety of online training packages. To list a few we have:

- Leap Frog Enterprise for early development
- Hooked on Phoenics for reading and spelling
- Houghton Mifflin Learning Technology for math

Assuming suitable business arrangements are negotiated, YED would be able to dynamically draw upon them as cloud functions in its interactions with the child.

Both parent and teachers would ideally provide the daemon with additional information about the child such as relevant medical information; teaching plans for the upcoming semester course grades, etc. Finally, note that external data could also come from other devices that the student interacts with. The use of robotic toys to engage children with autism has been shown to provide valuable information and to positively impact social and psychological development [30]. Such information would be extremely valuable to the daemon and given the increasing availability of high bandwidth access to the home and classroom, these robots could be considered extensions of the daemon.

2.3 Mentoring the Student

Once YED has adequately modeled a student, it can begin to execute its primary goal, i.e., develop the student's capabilities and enrich the student's life. In those areas where the student is not performing at the level of their potential, YED will determine the optimal training program. For example, if YED learns that the student has a problem with memorization, there are multiple games the daemon can play with the student that develop this skill implicitly.

Reading, writing are arithmetic are fundamental skills that can all be evaluated and enhanced via online applications [48, 47, 49].

In those areas where a child has a clear aptitude, the goal is to provide the child with opportunities to broaden and develop this talent. For example, if a child shows an interest in art, the daemon could arrange a virtual museum tour, all the while taking note of those artists that speak to this individual. If a child draws a picture while on-line, the daemon can compare it existing artworks and display like compositions.

The ability of a machine to learn concepts allows it to teach the concepts as well by inverting the steps. For each concept to be taught, an algorithm would generate an effective set of definitions, similar hypotheses with counterexamples and positive examples. In a variation of the game of "20 questions", the algorithm would interact with the student by providing successive instances from this concept set and allowing the child to converge on the underlying concept. The daemon would be able to verify that the child has indeed learned the concept using the same game. Computer aided instruction has been found effective in concept formation with even very young children [46].

In addition to helping the daemon learn about a student, games can be a powerful modality for teaching by bypassing the traditional one-way direction of the teacher student relation and facilitating implicit learning.

Chess, for example, has been shown to develop logical thinking, personal responsibility, and risk analysis. Furthermore, studies have shown that students who learn and play chess do better across the board. To illustrate the time series behavior of the YED note the following sequence used by the chess coach[2] in a Waterbury CT grade school chess club.

1. At age 6 the YED would introduce the child to the Fritz & Chester program. To test the child's level of proficiency, YED would periodically initiate a game with a standard chess playing algorithm and gauge the level of play.
2. At the appropriate time, say around age 10, YED would introduce the student to Chess Master periodically checking on the student's progress.
3. When the student reached the necessary level of maturity, he/she would be introduced to RYBKA or Fritz.

The use of games as mentoring tools extends to the social and emotional domains as well as the cognitive domain. Games can be used to develop: role playing and competition [24], focus [30] and collaboration [25]. Specific attention has been given to games that can improve the social and emotional development of children with autism. [21. 23, 28, 30] One of the goals of these games is to help autistic children learn to recognize

[2] We are indebted to Dave DaAngelis for many fruitful conversations on teaching Chess.

emotional states in their interactions with other human beings [23, 28].

Just as we require assurance that our daemon is capable of learning how the student thinks and functions, we must be confident that the daemon is actually accomplishing its goal of teaching the student to learn what he/she did not know beforehand or to perform better at a specific task. Although methods exist for assessing the impact and effectiveness of new technology with respect to remediation [22], additional tools will be required to vet the scope of this proposal.

2.3.1 Special Needs: The ongoing education of our children is not independent of their well being and many of our children have special conditions that can negatively affect their education if not managed very carefully. Diabetes and Asthma have become more prevalent in the population and can lead to very serious consequences if not managed tightly. At the same time the convergence of telecommunications and computing is delivering the capability to remotely monitor these types of conditions and intercede when appropriate. We can incorporate these capabilities naturally into our daemon and take advantage of the 2-way communications channels that link the daemon and child. It is important to note that these capabilities are already here and producing positive results. The last three decades has seen many applications of machine learning to diagnosis and treatment plans where algorithms emulate the clinical reasoning of health care professionals [43, 44]. In several recent pilot studies, asthmatic children that were provided virtual care achieved excellent therapeutic and disease control outcomes compared with those who received standardized office-based care only [40, 41]. Diabetic children achieved higher compliance in taking their medicine and tighter control of HbA1C levels [42].

By virtue of it knowledge and access to privileged information, a child's daemon process would be in an ideal position to provide another link in the child's lifeline.

1. The daemon can remind the child to take their meds on schedule and serve as another pair of virtual eyes on the scene. In the studies referenced above, texting was one of the main modalities used for this purpose and it is also one of the main communication links between child and its daemon. In particular this form of communication would be ideally suited for children with hearing impairments.

2. Since the daemon "lives in the cloud" it could have access to al remote monitoring data and by applying the appropriate classification algorithm to the data serve as a surrogate / aide for the child's primary care physician.

We have discussed how multimedia games in particular are a powerful pedagogical tool that can be employed by the daemon for both learning and mentoring. There is work in progress that would extend these benefits to those children that are visually impaired [45].

2.4 Ubiquitous Access

As with the fictional daemons, YED would always be available when needed. Since the daemon preserves state information, it would have to live in the clouds, i.e., it would exist as a function that could be accessed over any broadband connection to a PC, Smart phone, Smart TV, etc. Since the daemon is not resident in any specific device, a child could start a tutoring session at home initially using their TV or PC, continue the lesson in the bus on the way to school and complete the lesson later that day at the ice cream parlor without skipping a beat or having to start over. Furthermore, if given permission by the student, YED would be able to determine when the student came on-line and autonomously request a connection. The daemon might have discovered something of interest to its protégé and set a reminder to bring it to his/her attention.

Figure 1 provides a high level diagram showing the daemons as virtual entities "living" in the cloud. Specifically, these entities are intelligent applications running on virtual machines within a set of physically diverse internet data centers. These entities have access to all the resources we discussed above as well as to the internet and all the major wireless networks over high speed optical transport. The increasing adoption of smart phones and internet enabled tablets that will allow rich communication between the daemon and child necessarily require high bandwidth between the subscriber and the daemon network as well. Fortunately, in addition to computing power, Moore's law generalizes to storage and telecommunications. Edholm's law [50], states that wired, wireless and nomadic network capacity will converge in about two decades. Furthermore, all three show similar, though time displaced exponential increases in data rate. Evolving technologies such as WIMAX, LTE and Femtocell are already being deployed globally and will deliver higher bandwidth as well as greater footprint, both of which are required for "ubiquitous access". IPTV, i.e., digital TV service is delivered using the IP protocol over a packet-switched network infrastructure, provides a two-way communication path between the subscriber the content providers allowing for interactive services.

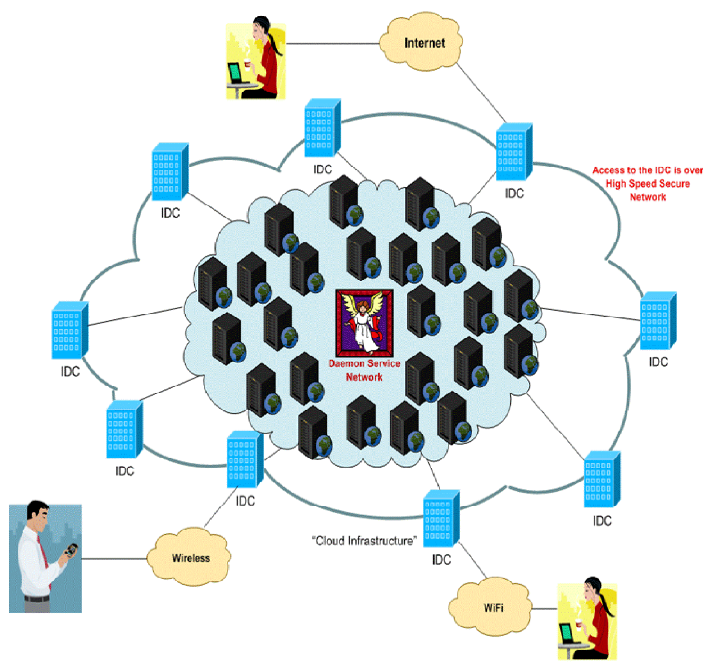

As we noted the daemon would be able to “converse” with its child “as if” it were a real entity. This allows the demon to evaluate the child’s verbal style and to provide feedback in a natural and encouraging manner. This is made possible by powerful developments in speech recognition and “text to/from speech” conversion. A noteworthy example of such technology is provided by the WATSON speech engine developed at AT&T Labs Research Organization. The following figure and description plus an extensive amount of additional information on related projects and more can be found at their website [51].

“WATSON is AT&T's speech and language engine that integrates a variety of speech technologies, including network-based, speaker-independent automatic speech recognition (ASR), Natural Voices® text-to-speech conversion, natural language understanding (which includes machine learning), and dialog management tasks”

Figure 2: AT&T WATSON Speech Technologies

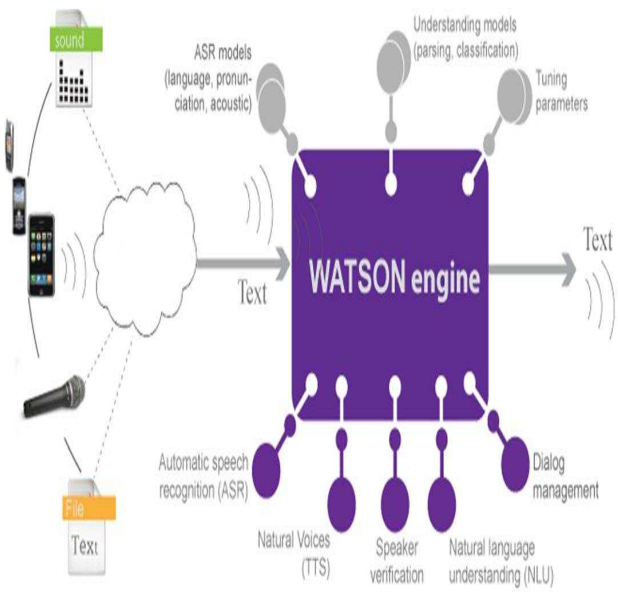

Chat rooms offer yet another communication pathway for sharing ideas, opinions and knowledge. A team of researchers built a dialogue system where users could chat with each other using natural language text and special verbs iodinating social gestures. The system was then extended using speech recognition technology to provide a 2-way path between users speaking over a phone and users accessing the system over the internet [54]. The relevant finding for our purpose, however, is that one of the most popular components of the systems is Cobot, a software agent that chats with the users and provides them statistical data about their chat sessions.

The communication between the daemon and its child would extend beyond speech to take advantage of all forms of multimedia and especially texting. The following chart compiled by the Nielsen service [52] shows how prevalent texting has becomes as a communication medium. In teenagers experience, texting is easier and faster. Multimedia Messaging Service (MMS) offers an additional attraction over SMS with the capability to combine sounds and images into the message. The obvious potential for educational purposes makes this an ideal modality for daemon – child interaction.

Figure 3– Trends in Mobile Adoption

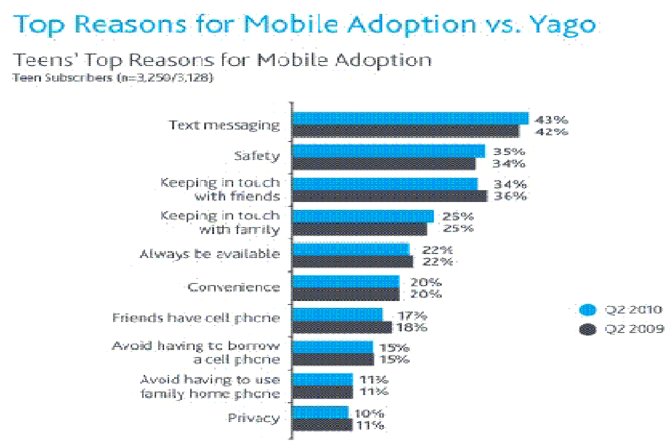

2.5 Putting the Pieces Together

We have discussed the individual components of the proposed daemon. The task at hand is to assemble the cross disciplinary team to build the application. The application must present itself as a coherent entity to the students so that the students will believe that YED understands them within the context of their educational life. If done correctly, YED will have passed the famous Turing test [16]. In the appendix we present 2 high level use cases of YED interacting with a student that illustrate the integrated functioning of all the components we discussed above.

3. CONCERNS, EXTENSIONS, AND NEXT STEPS

Recognizing that the very high level outline we have presented above is already quite ambitious, it would be unfortunate if we did not at least mention potential issues, extensions and next steps.

3.1 Concerns and Absolute Constraints

Because YED has access to and manipulates personal information about its students, there are clear security and privacy issues that must be addressed. The daemon must also be completely autonomous and not subject to the control of outside parties, no matter how noble their intentions.

We also do not view the YED as a substitute for the teacher student relation but as an adjunct. YED can provide both parents and teacher with status updates and logs of its interactions with the student. However, even here, great care has to be taken to protect the privacy of the student and the integrity of the relation between YED and the child. We are this lead to the following set of absolute constraints on the construction of the daemon.

Privacy:

- Personal data must not be able to read or accessed by any one outside the <human, daemon> pair or a legally approved surrogate.

Security:

- The set of individuals that can provide data to a daemon must be legally empowered to do so.
- Any data update must leave behind an audit trail that can be viewed by the daemon's charge or approved legal authorities.
- The mission / directive of a specific daemon can only be modified by the charge of that daemon or by legally approved surrogates.
- Inter-daemon communication is encrypted and shared information is not kept in permanent state.

Advocacy:

- A daemons goal set is completely focused on the advancement of its charge where the meaning of this evolves as the bond between the daemon and its human becomes progressively richer.
- The daemon must be bound to in no way encourage illicit or unlawful behavior on the part of its charge.
- The daemon must maintain a log of its dealing with its mentee and this log can only be made available to outside parties with clear legal authority

For the daemon to be effective in learning about the students it serves there must be significant and continuous interactions between both parties. This is one of the reasons that it must be always available, accessible over a multitude of media pathways and easy and pleasant to interact with. This will require a significant investment on the part of the student.

3.2 Extensions

In this short paper we have focused almost exclusively on the purely educational role of the daemon in its interactions with the student to which it is attached. The daemon, however, can serve multiple purposes and offer a wide variety of support to humanity. Here are a few.

1. It can reduce the possibility of the student falling behind by proactive monitoring of the subject's proficiency relative to the standard curriculum.
2. Assume the daemons are able to communicate with each other and share information about their charges' preferences, strengths and talents. This would allow them to create group events. This would allow them to create social interactions between likeminded individuals by opening communication and sharing of work and ideas. Chat rooms and internet facilitated conference calls can be set up either dynamically or by reservation.
3. It can provide the student an "always available" friendly playmate, where play and studying are interwoven.
4. It can provide an objective third party view of the student's strengths and weaknesses to both teachers and parents.

The second and third capabilities would be especially useful for students that are remote from urban centers or with disabilities that do not allow them access to other students on a regular basis.

3.3 Next Steps

In this paper, we propose that it is now feasible to develop an intelligent agent capable of learning the cognitive sand emotional structures of a child and which can then use that information to facilitate the child educational and social development. In summary, YED would be built to embody the following key characteristics.

1. The daemon possesses a deep understanding of the human to which it is attached
2. The daemon has access to a realm of knowledge no individual human being could attain.
3. The daemon is an advocate for its charge, i.e, it can only act in ways that positively affect the development of its charge
4. The daemon is always available
5. Evolves with its charge
6. Is personable, discrete, and secure

A fascinating experiment has occurred over the last decade that show children given access to computers with very little supervision have been able to learn how to use these machines to educate themselves [53]. The structure of the response indicates a self-organizing system at work. It is within our grasp to provide every child with access to the best our technology can provide. It is within our grasp to give every child access to an advocate that can grow with them and guide them to fulfill their innate potential.

To move from vision to reality we propose the next steps.

1. Assemble a team of educators, psychologists, computer scientists and data networking specialists to create the design specifications for the daemon.
2. Provide a proof of concept that an intelligent agent can learn the developmental level of a student and then provide the appropriate set of remedial exercises, games and training packages that address the individual's needs. To accomplish this in a timely fashion we would narrow the focus to grades 6-8.
3. Create the high level software design for the daemon paying special attention to ubiquitous access and the data network infrastructure that would be required.

4. USE CASES

NAME	**1.0 Math Homework**
REFERENCES	Harcourt Math Teacher Edition; Assessment Material
DESCRIPTION	This use case will identify grade level achievement, identify and correct pre-existing knowledge gap, and test corrective action
ACTORS	Student, YED
ASSUMPTIONS	Student in grade 6 Homework: Linear equation
PRE-CONDITIONS	1. Math Homework can be downloaded form schools website 2. Harcourt Math Teacher edition (or equivalent) is accessible to YED 3. Math assessment software is accessible to YED
BASIC COURSE	1. Student starts to work on homework assignment 2. Student connects to YED via home computer 3. Student logs into school/classroom site to download math homework 4. Student starts solving simple equations 5. YED monitors the process and identifies mistake in the final result due to recurring mistake in simplification process. 6. YED accesses the assessment software and identifies cause as issue with division process (prime divisors) 7. YED communicates to student that the results of the equations are not correct due to above issue 8. YED downloads interactive material to explain prime divisors and division process 9. YED offers examples and exercises for the student to work 10. Student works the exercises. 11. YED tests the application by suggesting news equation to be solved 12. Student solves them correctly 13. YED verifies that student has learned the concept by checking that exercises are solved h correctly 14. Session terminates
POST-CONDITIONS	State information is written to YED memory. User model is updated to show attainment of proficiency with prime divisors

NAME	**2.0 YED as Companion**
NOTES	This experience can be achieved on application of a Smartphone (mobile device) or desk top
DESCRIPTION	This use case will use multi disciplinary e-educational game to keep the student engaged and challenged
ACTORS	Student, YED
ASSUMPTIONS	1. Student in grade 7 2. YED has learned the following profile of the student a) Short attention span b) Competitive c) Successful "multitasker" d) Limited vocabulary for the grade
PRE-CONDITIONS	1. Adaptive e-educational games accessible to the YED 2. Assessment software is interconnected with e-game
BASIC COURSE	1. Student accesses YED for entertainment purposes 2. YED suggests an adventure based challenge game 3. YED uses a multi-disciplinary game including social studies, science, and math 4. Within each challenge, the YED includes neologisms, along with the definition 5. YED observes that the student excels in social studies and math and increases the difficulty level for those areas 6. YED observes that the science tasks are not accomplished properly, therefore it decomposes the macro science challenge into intermediate steps, 7. YED analyzes the results of the game, compares against assessment data to identify potential gaps 8. YED identifies remediation steps. 9. YED stores the final conclusions 10. Session terminates
POST-CONDITIONS	YED will use the stored data to propose new interactions with the student, updates student profile

5. REFERENCES

[1] Plato, **The Works of Plato**, **Translated by B. Jowett,** Tudor Publishing Company, New York, Vol. 2, pp 415, 1945.

[2] Adema, J., J., Models and Algorithms for the Construction of A Revised Simplex method for Test Constructions Problems. **Research Report 90-5, Department of Education University of Twente, Netherlands.** 1990.

[3] D. Angluim, Queries and Concept Leraning, **Machine Learning**, Vol. 2, 1985, pp. 219-342.

[4] D. Angluim, Learning Regular Sets from Queries and Counterexamples, **Information and Computing**, 1987, 87:87-106.

[5] C. Bauckhage & C. Thurau, Exploring the fascination: video games in machine learning research and education, **Proc. of the 2nd International Workshop in Computer Game Design and Technology**, 2003, pp. 61-70.

[6] E. M. El-Alfy & R.E. Abdel-Aal, Construction and Analysis of Educational Tests Using Abductuve Machine Learning, **Computers and Education**, Vol. 51, Issue 1, 2008, pp. 1-16.

[7] R. B. Fisher, A Review of Linear Programming and its Applications to the Assessment Tools for Teaching and Learning (asTTle) Projects., **Technical Report 5, Project asTTLE,** University of Auckland, 2000.

[8] A. Jennings & H. Higuchi, A User Model Neural Network for a Personal News Service, **User Modeling and User-Adapted Interaction**, 3910, 1993, pp. 1-25.

[9] A. Kobya & W. Wahlster, **User Models in Dialog Systems,**. New York:, Springer Verlag, 1990.

[10] F. M. Lord, **Applications of Item-Response To Practical Testing Problems**, New Jersey: Lawrence Erlbaum, 1980.

[11] T. Mitchell, **Machine Learning**, San Francisco: McGraw-Hill, 1997.

[12] T. Pack, Pandora let's users create online Radio Stations (Link-Up @Home). **Information Technology**, 2008.

[13] R. Rivest, Learning Decision Lists, **Machine Learning**, Vol. 2, No. 3, 1987, pp. 229-246.

[14] A. Torcher, R. Jahrer, R. Legenstein, Neighborhood-Based Algorithms for Large Scale Recommender Systems and the Netflix Prize Competition, **KDD 08**, ACM Press, 2008.

[15] K.T. Sun & S.F. Chen, A study of Applying Artificial Intelligence Techniques to Select Test Items, **Psychological Testing,** 46(1): pp 75-88.

[16] A. Turing, Computing Machinery and Intelligence, **Mind,** 59(236), 1950, pp. 439-460.

[17] L. G. Valiant, A Theory of the Learnable, **Communications of the ACM**, 27(11); 1984, pp.1134-1142.

[18] W. J. Van der Linden, **Linear Models for Optimal Test Design**, New York: Springer Verlag, 2005.

[19] Zhang & Byoung-Tak, Cognitive Learning and the Multimodal Memory Game – Toward Human-Level Machine Learning, **IEEE World Congress on Computational Intelligence**, 2008.

[20] J. Joyce, Pandora and the Music Genome Project, song structure analysis tools facilitate new music discovery. **Scientific Computing,** 2006.

[21] U. Anderson, P. Josefsson & L. Paereto, Challenges in designing virtual environments training social skills for children with autism, **Proc. of the 3nd International Conference on Disability, Virtual Reality and Assoc. Technology,** 2006, pp. 35-42.

[22] C. A. Wissick & J. E. Gardner, Conducting Assessments in Technology Needs: From Assessment to Implementation, **Assessments for Effective Intervention,** Vol. 33, No. 2, 2008, pp78-93.

[23] S. Barron-Cohen, Autism: The Empathizing-Systemizing (E-S) Theory, **The Year in Cognitive Neuroscience 2009: Ann. New York Academy of Sciences,** 1156: 2009, pp. 68-80.

[24] B. M. Slator & H. C. Chaput, Learning by Learning Roles: A virtual role-playing environment for tutoring, **Proceedings of the 3rd International Conference of Intelligent Tutoring Systems.**. Springer-Verlag, 1996, pp. 668-676.

[25] I. Martinez-Ortiz, P. Moreno-Ger, J. L. Sierra, et al., Production and Development of Educational Videogames as Assessable Learning Objects, **First European Conference on Technology Enhanced Learning,** 2006. pp. 316-330.

[26] E. Gouli, H. Kornilikas, K. Papanikolaou, et al., Adaptive Assessment Improving Interaction in an Educational Hypermedia System. **Proc. of the Pan-Hellenic Conference with International Participation in Human Computer Interaction,** 2001, pp. 217-222.

[27] B. Zupan, M. Bohanec, J. Demsar et.al. Learning by Discovering Concept Hierarchies, **Artificial Intelligence,** 109: 1999, pp. 211-242.

[28] R. E. Kaliouby, R. Picard & S. Baron-Cohen, Affective Computing and Autism, **Ann. New York Academy of Sciences,** 1093: 2006, pp. 228-248

[29] F. Fovet, Using Distance Learning Electronic Tools within the Class to Engage ADHD Students: a Key to Inclusion?, **3rd ASEE/IEEE Frontiers in Education Conference,** 2007.

[30] F. Michaud & C. Theberge-Turmel. Mobile Robotic Toys and Autism, **Socially Intelligent Agents – Creating Relationships with Computers and Robots**, ed. K. Dautenhahn, A. Bond, L. Canamero, et. a.l., Kluver Academic Press, 2002.

[31] C. McCormack & D. Jones, **Building a Web-based Education System,** Wiley Computer Publishing, 1997**.**

[32] C. S. Nave, R.A. Sherman, D.C. Funder, et. al., On the Contextual Independence of Personality: Teachers' Assessments Predict Directly Observed Behavior After Four Decades, **Social Psychological and Personality Science**, Vol.1 No. 4, 2010, pp 327-334.

[33] R. R. McCrae & P. T. Costa, in The Five-Factor Theory of Personality, in **Handbook of Personality: Theory and Research (3rd ed)**, O. P. John, R.W. Robins & L.A. Pervin (ed), New York, Guilford, 2008.

[34] D.P. McAdams, What do W Know When We Know a Person, **Journal Of Personality**, Vol. 63, 1995, pp. 363-396.

[35] B. S. Bloom & D.R. Krathwohl, **Taxonomy of Educational Objectives: Book 1, The Cognitive Domain,** New York, Longmans Green, 1956.

[36] H. Franco, L. Neumeyer, Y. Kim & O. Ronen Automatic Pronunciation Scoring for Language Instruction, **Proceedings of the 1997 IEEE International Conference on Acoustics, Speech and Signal Processing,** Vol. 2, 1997, pp. 1471 – 1474.

[37] T. Cincarek, R. Gruhn, C. Hacker, et. al., Automatic Pronunciation scoring of words and sentences independent from the non-native's first language, **Journal Of Computer Speech and Language,** Vol. 23, Isss.1, 2009, pp. 65 – 68.

[38] S. Petersen & M. Ostenndorf, A machine learning approach to reading level assessment, **Journal Of Computer Speech and Language**, Vol. 23, Isss.1, 2009, pp. 89 – 106.

[39] Holte, R. C. Very simple classification rules perform well on most commonly used datasets. **Machine Learning**, 11:63-91, 1993.

[40] D. S. Chan, C. W. Callahan, V. B. Hatch-Pigot et. al., Internet-Based Home Monitoring and Education of Children With Asthma Is Comparable to Ideal Office-Based Care: Results of a 1-Year Asthma In-Home Monitoring Trial, **Pediatrics,** Vol. 119., No. 3, 2007, pp. 569 – 578.

[41] C. Wainwright & R. Wootton, A review of Telemedicine and Asthma, **Disease Management and Health Outcomes**, Vol. 11, No. 9, 2003, pp 557-563.

[42] S. Krishna, S. A. Boren & E. A. Balas, Healthcare via Cell Phones: A systematic Review, **Telemedicine and e-Health**, Vol. 15, No. 3, 2009, pp. 231-240

[43] K. A. Horn, P. Compton, L. Lazarus & J. R. Quinlan, An Expert System for the Interpretation of Thyroid Assays in a Clinical Laboratory, **Australian Computer Journal**, Vol. 17, No. 1, 1985, pp. 7 - 11

[44] A. Persidis & A. Persidis, Medical Expert systems: An Overview, **Journal Of Management in Medicine**. Vol. 5, No. 3, 1991, pp. 27-34.

[45] A. Buaud, H. Svensson, D. Archambault & D. Burger, Multimedia Games for Visually Impaired Children, **Proceedings of the 8th International Conference on Computers Helping People with Special Needs**, 2002, pp. 173-180.

[46] H. S. Ayvaci & Y. Devecloglu, Computer-assisted Instruction to Teach Concepts in Pre-school Education, **Procedia – Social and Behavioral Sciences**, Vol. 2, Iss.2, 2010, pp. 2083 – 2087.

[47] R. Moxley, B. Warash, G. Coffman et. al, Writing Development using Computers in a Class of Three-Year-Olds, **Journal of Computing in Childhood Education**, Vol. 8, No. 2/3, 1997, pp. 133 -164.

[48] C. Arroyo, What is the Effect of Extensive Use of Computers on the Reading Achievement Scores of Seventh Grade Students, **ERIC Technical Report** ED353544, 1992.

[49] B. Burley, **Hands-On Computer Activites for Teaching Math: Grades 3-8**, Jossey Bass, 2004.

[50] S. Cherry, Edholm's Law of Bandwidth: telecommunications data rates are as predictable as Moore's law., **IEEE Spectrum,** 2004.

[51] AT&T Labs Research WATSON Speech Engine url: http://www.research.att.com/projects/WATSONASR

[52] **Nielsen Wire**, U.S. Teens Mobile Report: Calling Yesterday, Texting Today, Using Apps Tomorrow, Oct. 114, 2010.

[53] S. Mitra, et. al., Acquisition of Computer Literacy on Shared Public Computers: Children and the "Hole in the wall", **Australasian Journal of Educational Technology**, Vol. 23. No. 3, 2005, pp. 407-426.

[54] M. Kearns, c. Isbell, S. Singh et. al. CobotDS: a spoken dialogue system for chat, **Proceedings of the Eighteenth national conference on Artificial Intelligence,** 2002, pp. 425 – 430,