TOBY: Early Intervention in Autism through Technology

Svetha Venkatesh, Dinh Phung, Thi Duong
School of Information Technology
Deakin University, Australia
svetha.venkatesh,dinh.phung,thi.duong@deakin.edu.au

Stewart Greenhill, Brett Adams
Department of Computing
Curtin University, Australia
s.greenhill,b.adams@curtin.edu.au

ABSTRACT
We describe TOBY Playpad, an early intervention program for children with Autism Spectrum Disorder (ASD). TOBY teaches the teacher – the parent – during the crucial period following diagnosis, which often coincides with no access to formal therapy. We reflect on TOBY's evolution from table-top aid for flashcards to an iPad app covering a syllabus of 326 activities across 51 skills known to be deficient for ASD children, such imitation, joint attention and language. The design challenges unique to TOBY are the need to adapt to marked differences in each child's skills and rate of development (a trait of ASD) and teach parents unfamiliar concepts core to behavioural therapy, such as reinforcement, prompting, and fading. We report on three trials that successively decrease oversight and increase parental autonomy, and demonstrate clear evidence of learning. TOBY's uniquely intertwined Natural Environment Tasks are found to be effective for children and popular with parents.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords
Autism, early intervention, therapy, wait-list

INTRODUCTION
As computer scientists, we are familiar with the sense of satisfaction that accompanies an accepted paper or grant proposal. Much rarer is a sense of true delight at a finished project. We describe here one such unusual project–the development, from concept to field trial, of an early intervention program for childern with Autism, TOBY Playpad (tobyplaypad.com).

Our journey began in early 2009 when our friends' 2 year old son was diagnosed with Autism Spectrum Disorder (ASD), a neuro-developmental disorder that causes deficits in social interaction, communication, behaviours and interests. None of us had first-hand experience with Autism, despite its prevalence of 1 in 89 children. Alarmingly, the diagnosis came without clear guidelines for the parents about what to do next. The experts stressed the importance of early intervention, but our friends felt helpless in the face of a bewildering array of different therapies, while they waited to access formal therapy. When they finally began therapy, they discovered it required 20 hours one-on-one with their son, and another 20 hours of preparation, each week. The cost in dollars, anxiety, and stress, was indescribable.

Our goal was simple: to help parents during the stressful time between diagnosis and commencement of formal therapy, and during therapy itself. We wanted to empower them to deliver therapy at home. We set the following specific goals for what ultimately became TOBY Playpad:

- Deliver stimuli flexibly within a *rigorous* learning framework that uses best-practice techniques from behavioural therapy, including reinforcement, prompting, and measurable criteria for skill mastery and syllabus progression;
- Teach from a *multi-skill* syllabus spanning visual and auditory understanding, receptive and expressive language, and, critically, social skills and imitation;
- Deliver stimuli in *mixed environments*, on- and off-device, intergrated with this rigorous learning framework;
- Maximise the scope by making it *language independent*.

This paper presents a progression from our previous work [12]. However, beyond some reiteration of necessary background, it is markedly different from [12]. Previous work [12] is a systems paper targeting a system venue; it focuses on the system architecture, database structure, construction of stimuli, and detailed specifications of the algorithms that adapt the complexity of stimuli and measure mastery. In contrast, This paper focuses on user-centered design and iterative redesign. Outstanding distinct novelties and contributions from [12] include the following. a) *Syllabus Evolution:* we provide a discussion of the characteristics of Applied Behavior Analysis (ABA), its aptness to computer-mediated delivery, and the flash-card based model that TOBY seeks to replace and augment. The description of changes to the feedback model for NET and the division of responsibility between System and Parent for the different task types is new. This section also includes a more detailed breakdown of the four major skill areas in the syllabus. b) *Imitation Tasks:* this section is newly added which discusses how video modeling is used to teach various skills, how we take feedback from the parent, and communicate prompting requirements. It also describes the interface changes between prototype and final implementation (cf. Figure 2). c) *Natural Environment Training* (NET): our implementation, both in protocol and interface, changed substantially from [12]. These changes and their rationale are explained in details. d) *Lesson Planning:* this section discusses the rationale and effect of interfaces changes which

CHI 2013, April 27–May 2, 2013, Paris, France.

was not implemented previously. e) *Help Guide:* TOBY relies on the parent learning a number of concepts basic to ABA in order to deliver therapy, and the help system is a crucial part of learning process. In previous trials parents received a degree of instruction from therapists, but in the new final longitudinal trial they received no such help. The discussion of the provided documentation and videos is also newly added. f) *Experiments:* the results for Trial 2 include an evaluation of evidence of learning, and a comparison of NET/iPad use, both of which are not presented in [12]; Trial 3 is newly added which further substantiates previous trials in several respects, including a larger number of participants and was not supervised by therapists – parents received all training through the application. Trial results show evidences of learning, both on- and off-iPad. In particular, the most novel aspect of TOBY, its Natural Environment Tasks, were well received.

BACKGROUND

Software and Assistive Devices for ASD

Research into assistive devices for ASD have targeted three main areas: early diagnosis and progress measurement; affect recognition; and social skill development. Kientz et al. used toys fitted with accelerometers and wireless interfaces to trigger video recording for review and pediatric analysis [6]. Blocher and Picard developed mobile, assistive devices to help children recognize affect (a common difficulty for children with ASD) using facial analysis [2]. Exploratory work by Schmidt and Schmidt examined 3D environments used as surrogates for the real world in order to teach social skills [9], whereas Stanton et al. experimented with the effect of real-world, robotic toys on social development (i.e., autistic behaviours) [11].

Focussing on visual support, many systems have demonstrated efficacy in early treatment of mental disorders including autism (e.g., [1, 10]). Examples of proprietary software include *DTTrainer*[1] and the Picture Exchange Communication System (PECS) [3]. Video has also been used, in the form of social stories, which are concrete, idiosyncratic video narratives that teach social skills [7]. Computer-assisted intervention has proven effective in teaching language, reducing inappropriate verbalisation, and improving functional communication and generalisation [4, 5]. *Teachtown*[2] is perhaps the closest system to TOBY currently available in that it provides graded online and paper-based lessons rooted in Applied Behavior Analysis (ABA) theory, but it suffers from a restricted set of stimuli and impoverished adaptation to response.

Anecdotal evidence suggests the iPad has been a great learning tool for children with autism.[3] There is a host of iPad apps targetting specific skill deficits. E.g., *Proloquo2Go*, *iComm*, and *TapToTalk* for communication via symbols and text-to-voice; *Grace* for sentence-building via images; *iCommunicate* for storyboarding with pictures and images; *First-Then-Visual-Schedule* for daily schedules; *AutismExpress* for emotion interpretation; and *Stories2learn* for personalized social stories about complex social situations. But none of these systems target systematic, automated stimulus generation and complexity adaptation, at the level proposed here.

[1]www.dttrainer.com
[2]www.teachtown.com
[3]http://www.sfweekly.com/2010-08-11/news/ihelp-for-autism/%20./

Behavioural Early Intervention Therapy for ASD

Applied Behavior Analysis (ABA) [8] and Discrete Trial Training (DTT) are prominent established therapies (US National Standards Report on Autism, 2009). They focus on cognitive functions, such as object labelling and categorization. While intervention is designed and partially administered by therapists, much of the therapy requires parental involvement. ABA's reliance on paper-based materials can translate to significant preparation costs, and suffers from difficulty in communicating dynamic concepts, such as verbs. ABA and DTT also require fine-grained record keeping during sessions, which in practice can be poorly implemented.

But ABA's structured progression formulae, record-keeping, and stimulus-based learning, are a good fit for computational delivery. Thus ABA underpins the content delivery framework described below, and instantiated as TOBY Playpad. But we hasten to add that this framework is not a replacement for human-delivered therapy, but a complement to it.

SYLLABUS EVOLUTION

Our initial vision of TOBY Playpad arose from direct experience with ABA therapies. TOBY aimed to simplify delivery of these therapies for parents, by automating the tedious or time-consuming aspects, including: making flashcards, printing, and laminating; sharing images between parents; and keeping records of therapy sessions. In 2009 we designed a multi-touch surface to replace table-top activities done using hand-made flashcards. This prototype stored and presented images and sounds for matching tasks, and recorded performance. But by April 2010, the iPad offered a dramatic improvement in price and robustness of multi-touch devices.

At that time we contacted a support group for parents of children with autism, Autism West. Together we formulated the goal of developing a product for early intervention, to cover the crucial time between diagnosis and access to formal therapy, which could be up to six months. Autism West supplied therapists and psychologists, who brought their own perspective to the problem. The result was a dramatic increase in the scope of skills covered by the syllabus, which grew from sensory and language matching tasks to include social, imitation, and Natural Environment Tasks.

We decided to use ABA for our therapy framework. ABA uses operant conditioning to teach a child how to respond to stimuli. While it can be used to teach a vast array of skills, for children with autism we are also interested in communication skills. These are founded on a range of pre-requisite cognitive and social skills. We began with ABA's use of flashcards to teach children object names, which is typical of computer-based autism therapy. But many skills must be taught in parallel. E.g., making eye contact; attending to another person; understanding and following basic instructions; focusing attention on a task; imitating gross- and fine-motor actions; and producing basic speech sounds.

A computer cannot teach a child to do these things, but a computer can *teach a parent* to teach a child. Thus our syllabus has been designed from the outset to involve both parent and child. TOBY teaches parents the basics of operant conditioning, which they then apply to the crucial areas of their child's development that cannot be handled by the computer.

In an ABA therapy model, we present a *stimulus* to which the child responds. If the *response* is the desired response, we give *reinforcement*–e.g., verbal praise, a favourite food or toy. If the response is not the desired response, we may *prompt* the child by demonstrating the correct response. Prompting is *faded* over time so that eventually responses are self-initiated (and similarly for reinforcement). *Measurement* of behaviour is used to decide: if a response is correct; when to reinforce; when to prompt and how much; when a skill has been mastered; and what should be taught next. Underlying the therapy is an algorithm that adjusts stimuli, reinforcement, and prompting as a result of responses. For those situations where the computer can measure a response directly, it delivers reinforcement and prompting, e.g., where the task requires the child to find a given stimulus picture from among a set of pictures presented on the screen. We call this type of task a *Solo* task, because it can ostensibly be performed by the child without assistance. Solo tasks are important, but they only cover a small subset of skills we want to teach. The remaining skills involve the parent in some way, either to recognise the child's response (e.g., an action or spoken word), or to present stimuli, prompts, and reinforcement. We call these *Partner* tasks, because the parent and child work must together.

The iPad delivers two kinds of language task: *receptive* and *expressive*. For receptive language, the child must learn to recognise an object given its name, and point to the correct picture in response to a voice prompt, such as "Find apple". For expressive language, the child must say an object's name given a picture and voice prompt, such as "What is this?" Expressive tasks are Partner tasks, as the parent must indicate if the response is correct, and prompt if necessary. While receptive language usually develops faster than expressive, TOBY treats these skills as independent, as children with autism often do not follow expected developmental paths.

Once the important concepts were identified, we collected images which were Creative-Commons licensed for use as stimuli in sensory and language tasks. We also recorded videos for imitation tasks, which extensively use video modelling.

Therapists stressed the importance of integrating computer-based therapy with real-world activities. Skills must be transferred to natural settings so they are generalized, and the child's reliance on screen-based prompting and reinforcement is reduced. TOBY's *Natural Environment Training* (NET), which is performed off-device, serves this purpose.

We experimented with several models for NET. Initially we tightly coupled the real and cyber worlds, by using digital exemplars to guide physical activities (e.g., an image of a banana, to be found while shopping). This model required parents to input many trial-level outcomes, so that the syllabus progression logic could be as sensitive as it was for iPad tasks, but this proved to be impractical and restrictive.

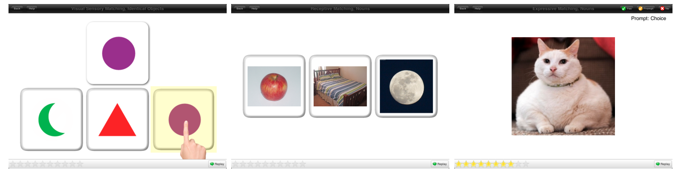

Figure 1. Matching tasks. Left: Sensory matching – the instruction is "Find same". An incorrect response causes TOBY to indicate the correct response using a pointing finger. Middle: Receptive matching – the instruction is "Find apple". Right: Expressive matching – the instruction is "What is this?". The child must say the word "cat". Parent inputs feedback using the buttons at the top right of the screen.

We switched to a model requiring a single response after each NET activity. This allows parents independence from the iPad while performing the task, but obtains the feeback necessary to drive syllabus progression, albeit at a larger granularity than the iPad tasks. Each NET instructs the parent how to perform the task, and how to prompt and reinforce if required. Every Solo and Partner task has a corresponding NET task.

	Task Type		
Role	Solo (iPad)	Partner (iPad)	NET
Stimulus	System	System	Parent
Response	System	Parent	Parent
Prompting	System	Parent	Parent
Reinforcement	System	System	Parent
Adaptation	System	System	System

Table 1. Roles of parents and system in Solo, Partner and NET activities.

The differences between the task types–Solo, Partner and NET–in terms of which agent performs the different parts of therapy is shown in Table 1. Observe that in all cases the system decides what lessons should be offered based on the performance of the child.

Our syllabus is divided into four major skill areas:

- ***Sensory***: perception and discrimination of sensory cues–e.g., colour, shape, same-ness and difference
- ***Imitation***: copy an action, design, or pre-speech sounds
- ***Language***: recognition and production of object names
- ***Social***: inter-presonal skills, such as joint-attention

There are between 10 and 15 specific skills in each of these areas, and a total of 51 skills in the syllabus, which are strucutred as a graph of skill dependencies. New skill nodes unlock as pre-requisites are completed. Each node may have multiple tasks, and all nodes have iPad tasks, except Social nodes, which only have NET tasks. The syllabus contains a total of 326 tasks: 34 iPad and 292 NET.

USER INTERFACE

Here we outline the different parts of the TOBY Playpad app from a user's perspective. Figure 2 is a simple schematic of the navigation choices in the app. The parent begins by registering a new account, which enables their profile and activity data to be stored for backup, migration across devices, or sharing with a therapist. If already registered, they login with their private credentials. From the *manage* screen, the parent can add, remove and edit child profiles. After selecting a profile, TOBY displays the *plan* screen, which displays a list of available tasks and an overview of the child's progress through the syllabus. The *guide* explains through text and video how to set up and perform different types of syllabus activities, including knowledge necessary to perform each task, such as how to prompt, reinforce, and motivate the child.

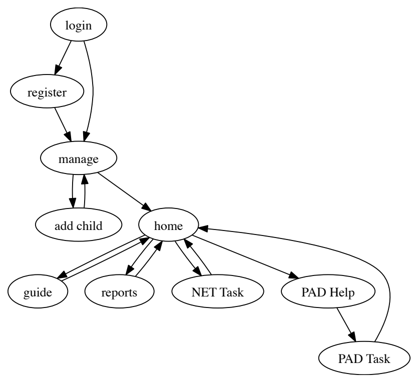

Figure 2. Simple schematic of the navigation choices in TOBY Playpad.

When the parent selects a task for their child to perform, and before proceeding to the task screen, they are first shown any task-specific instructions. Recall that there are three types of task in the system: *solo tasks*, which include receptive matching games the child can do under parental supervision; and *partner tasks*, which include imitation tasks, where the child must copy an action shown in a video, and expressive matching tasks, where the child must name objects presented on the screen; and NET tasks, which are performed off-iPad. In partner tasks, the parent provides feedback to TOBY about the child's response, and TOBY indicates how the parent should prompt the child if the response is incorrect.

The following sections describe the design of the different activities comprising TOBY's syllabus, including matching tasks, imitation tasks, NET tasks, as well as how lesson plans are communicated to the parent.

Matching Tasks

Matching tasks are used to teach categorisation. The child must identify objects that are similar in sensory or semantic properties. Recall that TOBY uses multiple exemplar training (MET), in which three categories are taught simultaneously, with the target category changing from trial to trial, and the non-target categories serving as distractors.

The simplest kind of matching is sensory. Here TOBY presents a target image, and a row of alternatives, with the instruction, "Find same". One of the images matches the target, and the other two are distractors, and the child must select the correct image by touching it, or dragging the target to it. Figure 1 (left) is a screenshot of a sensory matching task. The target (top) is the purple circle, which the child must pick from the bottom row of three images. This figure also shows one of the ways TOBY prompts the child for the correct response when an incorrect answer is given, in this

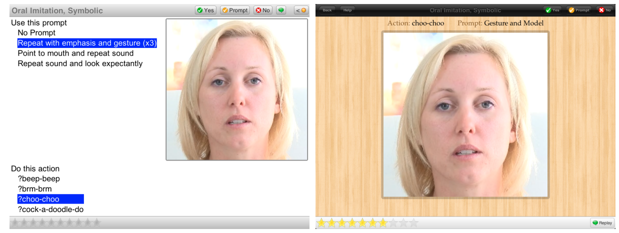

Figure 3. Left: Early prototype imitation task. Video models an action the child must imitate. The parent records their child's response with the buttons at the top right of the screen. Right: Subsequent redesign, with lists for prompt type and action removed to avoid mis-interpretation as choices.

case by highlighting the correct alternative, pulsing its size, and pointing to it. Once prompted, the child is allowed to choose again. A more complex version of this task presents two non-identical examples of an object. This trains the child to attend to similarities, whilst ignoring irrelevant details.

In addition to prompting after incorrect answers, TOBY offers rewards for correct answers. Between individual trials, reinforcement is given using a display of fireworks and verbal praise. Each correct answer earns the child a star, and 10 stars earns one token that can be traded for 20 seconds of play time on one of the built-in reward activities (e.g., bubble pop, drawing, favourite video, spinner). A progress bar at the bottom left of screen shows elapsed time, and after 20 seconds TOBY pauses the reward activity and returns to the current task.

Matching is also used to teach the names for classes of objects. Again, three images are presented, but now the target is defined via a voice instruction, such as "Find apple". Figure 1 (middle) shows an example of this task. TOBY organises categories into groups, and ranks them in order of increasing complexity. In this example, it has randomly chosen three sub-categories of *noun*: food, household object, and outdoor object, and has chosen the next simplest category in each set which has not already been presented: apple, bed, and moon. There are multiple examples of each category in the database, and these are chosen randomly during each trial.

The matching task described above that requires the child to choose the correct image given its name is termed *receptive matching*. A more difficult skill is to vocalize object names, which is termed *expressive matching*. Figure 1 (right) is an example of expressive matching. Again, TOBY uses three category MET, but in this case displays only the target image, with the instruction, "What is this?" TOBY relies on the parent to determine if the child's answer is correct, hence this is a partner task. Partner tasks include controls for the parent to give TOBY feedback. Here, valid feedback is one of Yes, Prompt, and No, which are mapped to buttons at the top right of screen. These are deliberately subtle in appearance so as not to attract the attention of the child.

If the child's response is correct, the parent presses the Yes button, and TOBY presents reinforcement before proceeding to the next trial. If the child's response is incorrect, the parent presses the No button. If appropriate, the parent may be asked to prompt the child using the prompt type displayed on the screen. In this example, the prompt type is *Choice*, so they would simplify the task by giving the child a choice, such as "Is it a dog or a cat?" If the child responds correctly with prompting, the parent presses the Prompt button.

A variety of tasks can be handled within this framework: Matching identical objects, where the child must match an object given an identical object; Matching non-identical objects, where the child must match an object to another of the same category; Receptive matching, where the child must match an object given its category; Expressive matching, where the child must name an object; Relational matching, where the child must match two objects based on a relationship between them. E.g., "Where does this [bird] live?" (in a nest), or "What is this [fork] used for?" (eating); and "What does not belong?" where the child must identify the object that does not belong in a group of similar objects.

Imitation Tasks

For imitation tasks, stimuli are in the form of video models, and the child performs the task by copying the actions of the model. Imitation tasks are used to teach skills in gross motor, fine motor, and oral motor imitation. These are prerequisite skills for expressive language. Each task includes 4 actions, which are presented in random order. According to performance, TOBY adjusts the suggested prompting level. Prompts are delivered by the parent using the suggested cues. TOBY includes 174 imitation videos, consisting of 100 actions and 74 echoic (oral-motor) videos. These are typically short, with an average duration of 1 second for echoic videos, and 4 seconds for action videos.

Figure 3 is a prototype screen for an imitation task, and a similar task in the current design. The task *symbolic oral imitation* requires the child to copy play sounds, like "beep-beep" and "meow". Most of the screen is reserved for the video, and above is the name of the action and current prompt type. Three buttons for parent feedback are located top right. This example has a prompt level of *Gesture and Model*, which the parent can enact to simplify the task by pointing to their lips, repeating the sounds, and breaking them into syllables.

The alterations from the early prototype, Figure 3 (left), were prompted by parent feedback. Initially, the system highlighted the current prompt level from a list of levels, and similarly for the current action. The aim was to communicate to parents the spectrum of prompt levels and action types by providing them in context. But in trials we found the effect of this interface was to confuse the parent, as it was interpreted as offering a choice, where in fact there was none. This screen also shows the parent controls rendered as regular buttons, which in practice attracted the attention of some children, who tried to press them and became frustrated when they were not allowed to do so. Subsequent versions were altered to the more subtle look shown on the right.

Natural Environment Training

Natural Environment Training (NET) generalizes skills by applying them to real-world situations. Typically, each part of the syllabus has 5 to 10 different NET tasks, grouped into two difficulty levels: *Adaptive* tasks modify a regular daily routine, such as meal time, dressing, washing, etc.; *Play* tasks are specifically designed games and play, which can be done at any place or time. The child must perform a mixture of adaptive and play tasks to demonstrate mastery of a skill. The protocol defined by the therapists is as follows: 1) *First*, the child must complete one adaptive task; 2) *Next*, the child must complete one play task; 3) *Finally*, the child must complete another adaptive and play task within 1 week (the only time constraint in the app) to demonstrate both skills have been retained.

This constraint complicates the interface design, since depending on the stage (1, 2, or 3) different tasks are available to the user. If the child fails one stage, they regress to the previous stage which must be repeated. Thus, the user interface must indicate which tasks are available and why, as well as the current stage within the sequence.

NET tasks have the following template: Goal; Materials required; Description; Instructions for prompting and reinforcement if required. TOBY allows the parent to navigate each NET task at increasing levels of detail via the Overview, Activity and Outcome screens. Figure 4 (left) shows the Overview screen, which lists the NET activities available to the child given their current stage. At the top of the screen is a summary of what is required for this stage (e.g., stage 3's summary is "Revision: Complete one Adaptive and one Play activity within one week.") The parent taps a task to see its goal and required materials. For more detail, the parent taps *Select Activity*, which displays a screen like Figure 4 (middle). The Activity screen has a full step-wise description of the activity, and instructions for prompting and reinforcement. It can be used as a reference when performing the task.

After the task is complete, the parent selects *Add Outcome* to provide feedback to TOBY about how the child performed, which displays a screen like Figure 4 (right). The Outcomes screen enables the parent to indicate successful completion, failure, or a range of outcomes corresponding to the three prompting levels. The parent selects the level that best describes their child's performance. After trials we found some parents chose not to perform tasks they believed their children could do, which prompted us to add the outcome: "Skip, my child can already do this." This is treated as a successful attempt, but records that the task was not performed.

The NET interface underwent several iterations. The prototype integrated browsing and feedback into a single screen (see Figure 5), and listed available tasks on the left (similar to the current version), but the right pane included *all* of the task's instructions. This was found to be overwhelming, and was altered to the multi-stage format described above to enable the parent to consume the information in bites. In the first prototypes, the completion state of each task was indicated by an icon (tick, cross) prefixed to its item in the list. Parents added feedback about a task by tapping this icon, which triggered a pop-up containing the response options. But some parents mistakenly interpreted this to mean they couldn't attempt a task more than once. That is, they assumed a ticked task was completed and no longer available, whereas repeated task attempts were in fact desirable. Subsequent designs removed all feedback about previous task attempts.

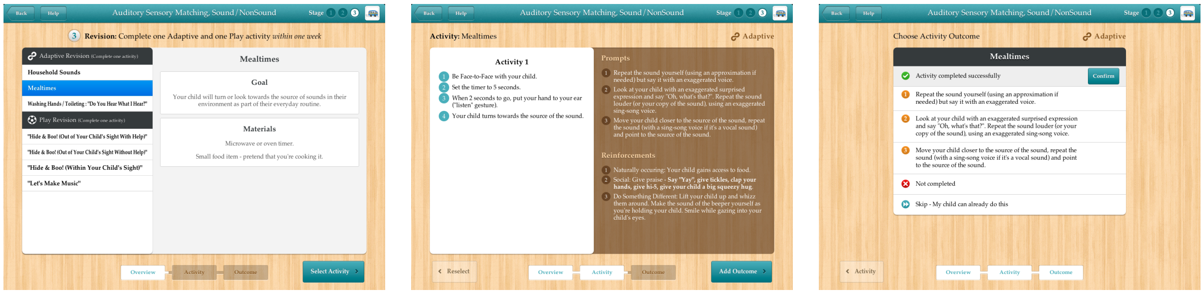

Figure 4. An example Natural Environment Training task: Overview, Activity, and Outcome screens.

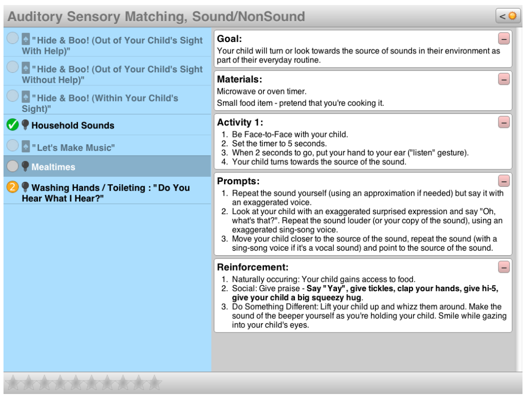

Figure 5. Early prototype NET screen.

Lesson Planning

Parents initiate tasks from the *Plan* screen, shown in Figure 6 (right). It is here that TOBY and the parent combine to schedule each day's lessons. The left side of the screen contains an extract of the current child's profile, including their name and photo (or an image of their choice), along with a summary of their progress. Bar graphs (bottom left) indicate the amount of the available course content completed for each kind of task–Sensory, Imitation, Language, and Social, in both iPad and NET–and so can be used by the parent to shape the lesson plan, by focussing on weak areas or capitalizing on sucesses.

On the right is the list of all currently available tasks. Each is represented by an item embellished with an icon indicating its type (left), its title (middle), and a selection toggle (right). Membership of this list is a function of the child's progress (recall that the syllabus is structured as a graph of task completion dependencies). The parent can tailor this plan by selecting a subset of tasks using the toggle "**+**", and those selected appear under the *My Plan* tab (middle top). My Plan was introduced to enable parents to firstly filter the–at times long–list of available tasks, and secondly to record tasks they intended to do. Tapping a task launches it, with optional intervening instructions for the chosen task. When it has been completed it appears for the remainder of the day checked with a tick, and is then moved to a list accessed under the *Completed* tab. Completed tasks can be revisited, but have no impact on progress. New tasks unlocked by the completion of pre-requisites appear dynamically in the available list.

The current form of the Plan screen is the result of a number of iterations that served to shift some of the initiative for lesson planning and pacing from TOBY to the parent. Figure 6 (left) shows an early prototype of the plan screen, then called "Today's Lesson". It contained only a list of available tasks and a legend. In this prototype, the task list remained static during each day (apart from marking completed tasks), and was refreshed only when the parent manually triggered a *Replan* with the button in the top right of screen. This "batch" protocol was adopted initially for two reasons: first, to render newly available tasks more noticable to the user, and, second, to cause the available tasks to be balanced across the syllabus.

Feedback from the first two trials caused us to rethink this protocol. While some parents became concerned they weren't doing enough tasks each day, others found having to wait until the next day to access tasks at which their children were excelling too constraining–they, and our therapists, were keen to capitalize on a child's progress in a branch of the syllabus. The solution was to make the list dynamically update, provide detailed reports on per-skill progress, and allow sub-selection of tasks (i.e., My Plan). Thus parents could perform tasks at their own pace, go as deep in any branch of the syllabus as they desired (subject to the completion of pre-requisites), and be informed about any imbalances in skill acquisitions so they could redress them in their own time.

Activity Reports

TOBY's range of reports grew as it was recognized that feedback about a child's progress was important to the parent, both for the encouragement it could provide, and as a guide to scheduling future lessons. Moreover, when TOBY is functioning as a "wait-list" support tool, fine-grained reports become valuable information about the child's development at the comencement of formal therapy. Visible in the lower-left corner of Figure 6 (right) is a miniature version of the course progress summary. The bars represent the proportion of available skills that have been mastered, grouped by task type: Imitation, Language, Sensory, and Social, for iPad and NET tasks. Sections shaded green are completed, cross-hatched sections are available, and unshaded portions represent course material yet to be unlocked. Depending on the situation, the parent may continue to focus on areas of strength, or may decide to focus on areas of weakness.

Other reports include: Progress achieved the previous day; Task attempts over time, colour-coded to indicate successful, prompted, and unsuccessful attempts–see Figure 8 (left); Task attempts for NET and iPad tasks; Daily summary of results for NET and iPad tasks–see Figure 8 (right).

Help Guide

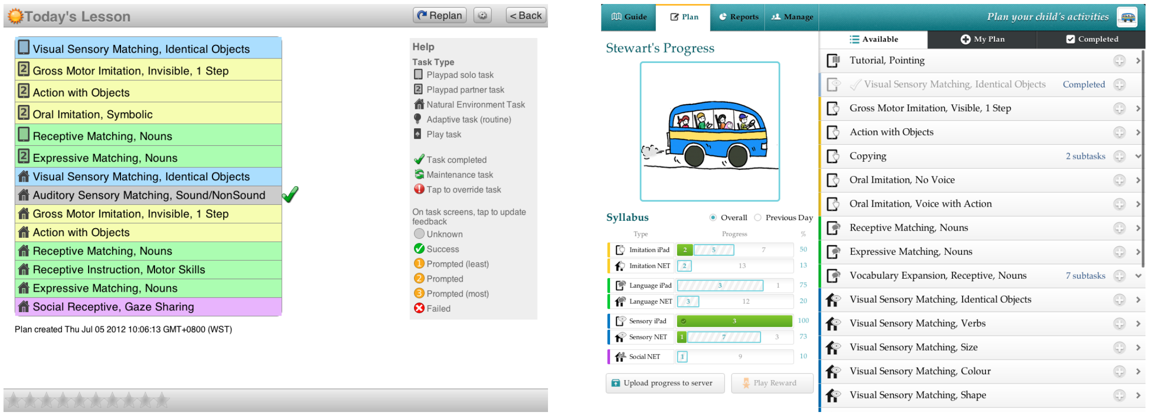

Figure 6. Left: Early prototype lesson plan screen (previously "Today's Lesson"). Right: Subsequent redesign showing a tabbed breakdown of available, scheduled, and completed tasks, and progress summarys.

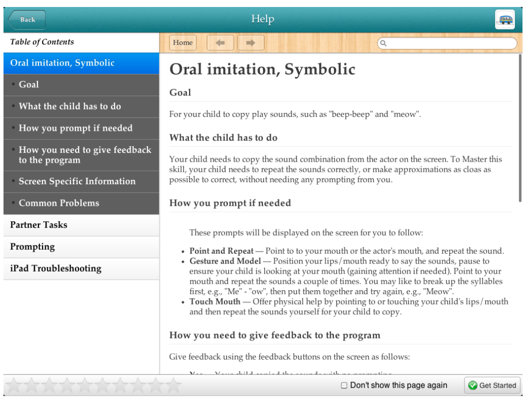

Figure 7. Detailed instructions specific to a task.

Feedback from the first two trials led to the creation of an extensive help system. Comprehensive help is unusual for iPad apps, where functionality tends to be limited, obvious, or inductive. In contrast, TOBY aims to empower parents in their role of therapist to their child, a task requiring them to digest concepts from the everyday language of psychology. Help includes: Information about early intervention, and *how* and *what* TOBY teaches; Instructions for Solo, Partner and NET tasks, prompting and reinforcement, and trouble-shooting about attention and motivation; Video of prompting strategies for a variety of tasks; and Specific instructions for every iPad and NET task. In total, the documentation includes 30,000 words, 31 minutes of video about the system, and 23 minutes of prompting examples.

Each task has a set of instructions which are presented whenever the task is started. Figure 7 shows the instruction screen for the task "symbolic oral imitation". It briefly describes the task, but importantly reminds the parent:

- How they should prompt if required
- How they should give feedback to TOBY
- How they can deal with common problems (e.g., their child can't perform the task, or isn't motivated).

After reading the information, the parent launches the task with the *Get Started* button at the bottom right of screen. As the parent becomes more familiar with TOBY, they can skip task instructions with the "Don't show this page again" checkbox. The instructions can later be accessed from within the task screen using the context-sensitive *Help* button. The rest of the documentation is available from the main screen under the *Guide* tab. Both task-specific help and general help share the same interface, shown in Figure 7. A topic list is shown on the left, and when a page is selected it expands to show the headings for quick navigation. Text is shown on the right, and may include supporting images and video.

EXPERIMENTS

We have undertaken three trials during the development of TOBY Playpad, each granting a greater degree of autonomy (and correspondingly lesser oversight) to the participants:

Trial 1: was conducted over 2 weeks, and involved 8 children and their parents, who performed tasks at Autism West, a registered provider of support services. This trial focussed on the mechanics of receptive and expressive matching, and an early implementation of NET. Each child spent 20-30 minutes on receptive or expressive matching (depending on developmental level) and the same amount on NET. Limited instructions about TOBY were given, and a therapist observed each session. Participants were solicited through Autism West.

Trial 2: lasted 4 weeks, and involved 16 children and their parents. This trial occurred in the homes of children. A parent information evening was conducted at Autism West to introduce TOBY's learning philosophy and interface. Two therapists called the families each week, and recorded any qualitative feedback in addition to TOBY's quantitative recording. Following the trial, a parent evening was held to de-brief and solicit further feedback or suggestions for improving TOBY. Participants were solicited through Autism West, and by expression of interest in response to an invitation on the TOBY website, up to a maximum of 16, first come first served.

Trial 3: lasted 6 weeks, and involved 47 children and their parents. Parents downloaded TOBY and worked at home at their own pace, and without therapist intervention, in a situation matching that faced by parents once TOBY was submitted to the App Store. Participants were solicited through Autism West and by expression of interest in response to an invitation on the TOBY website.

All children involved in the trials had a diagnosis of Autism, ranging in age from 2 to 8, with most between 2 and 6. There

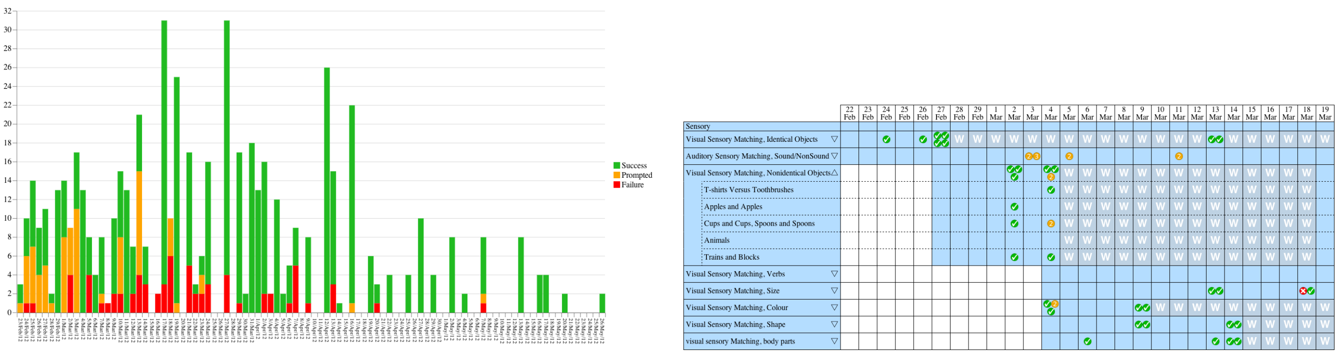

Figure 8. Left: Task attempt report, plotting successful (green), prompted (orange), and unsuccessful (red) task attempts per day. Ideally, reliance on prompting should decrease over time, as is the case here. Right: NET sessions report, showing times and results for each NET node, and optionally the related NET tasks. Note: a white background indicates a task has yet to be unlocked.

were approximately three times as many boys as girls, which corresponds roughly to the prevalence of ASD in the community at large. Below we focus on three aspects of these trials: qualitative feedback that informed design iterations; progression through the syllabus and evidence of learning; and profiling of when and how the application was used.

Qualitative Feedback, all Trials

Trials 1 & 2 were structured to obtain direct qualitative feedback from parents, and many has raised the impact of the subsequent design of TOBY and/or the suitability criteria for the application. Positive feedback is summarized under the following categories (paraphrased from explicit feedback):

A Sense of Empowerment – Many parents expressed their regret at not having had TOBY at the time of diagnosis, as they wanted to do something helpful but didn't know where to start. TOBY gave them easy, practical tools and ideas of how to do activities within the home, at any time. Some parents said TOBY taught them how to talk "therapist", and be proactive (based on the recorded progress) with their therapist, which made formal therapy time more productive.

Individuality & Adaptability – Parents liked how TOBY adapted to their child's progress at a fine level.

Syllabus Scope & Quality – Parents felt TOBY balanced play and structured activities. NET tasks were obviously designed by experienced therapists. Children displayed transfer of skills to other activities, such as being able to interact with computers generally. The syllabus helped identify gaps in current intervention programs and provide ideas on how to take therapy further, and served as a reminder for parents of different skills to focus on.

Recording & Accountability – Parents liked being accountable to TOBY and more thoughtful about therapy. Progress reports encouraged parents, and showed that previously taught skills have been maintained, strengthened, and generalised. The tangible record enabled them to say, "I've done enough today," rather than always, "I didn't get to X, Y, Z..."

Enjoyment & Ease of Use – NET tasks were easy and enjoyable, and parents look forward to extensions. Siblings joined in with tasks and activities–both NET and iPad–which provided futher reinforcement and family bonding.

Markee Moments – TOBY occasioned one child's first imitation of Mum, another's first pointing gesture, and another's first verbalization of "Mum".

Feedback obtained through exit surveys, communication with our therapists, and the parents events, prompted re-design of a number of facets of TOBY. A selection of these have been mentioned where appropriate, and are summarized here:

Algorithm – Some children weren't able to point, but would grasp the iPad, so we added tutorials to the root of the syllabus tree for the skills of pointing and dragging. More advanced children sometimes found progress tedious, so we introduced a "fast mastery" algorithm that rewarded 4 of 5 successful task attempts with completion (rather than 8 of 10).

Look & Feel – Some children were attracted to any and all buttons on screen. In response, parent controls on screens for children were de-emphasized. Some parents reported the initial NET screens to hold an intimidating amount of information. These were subsequently unpacked into parts, and combined with a visual indicator of the complex staging model.

Supporting Resources – Parents of newly diagnosed children were unfamiliar with core therapy concepts, such as reinforcement and prompting. The help guide was much expanded to cover these topics in depth and provide context-sensitive help at and during task performance. We also found that in some cases the home environment was interfering with therapy. In response we added material help on activity schedules, physical preparation of the home environment, use of external reward systems (e.g., to help with motivation), and explanation of learning concepts (goals, socially based activities, how to teach children).

Evidence of Learning, Trial 2

Based on iPad task performance, nearly all participants in Trial 2 can be assigned to one of 3 groups that showed evidence of learning. Figure 9 contains a plot for a child drawn from each group, where each plot presents the number of successful task attempts (S), successful attempts with a high. medium, and low levels of prompting (HP, MP, LP), and failed attempts (F). Below we interpret each group.

Children in Group 1 (profile ids 35, 37, 40, 50, and 72) progressed through the syllabus with ease, evidenced by consistent successful task attempts. Towards the second half of the trial when tasks get more challenging, the participants in this

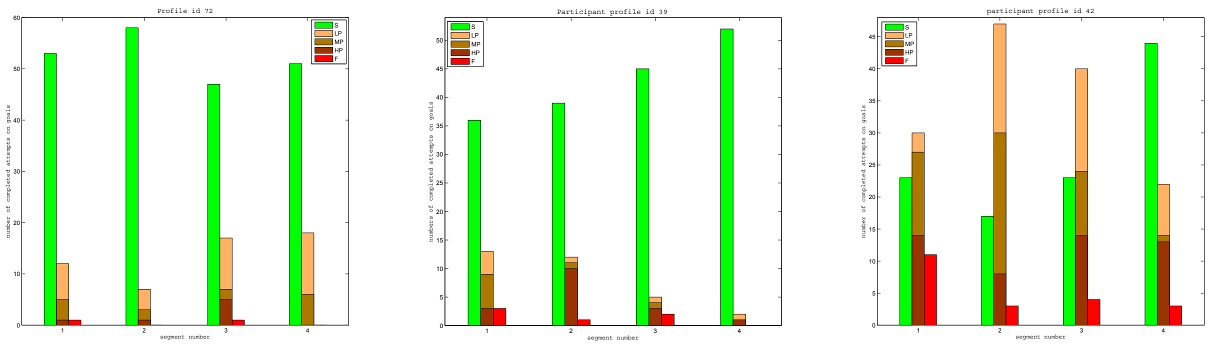

Figure 9. Trial 2 Evidence of Learning. Left: Group 1, successful task attempts over duration of trial; Middle: Group 2, steady increase in successful task attempts; Right: Group 3, increase in successful task attempts with decreasing prompts.

group typically learned new skills at a lower level of prompting. Figure 9 (left) plots an exemplar from this group.

Children in Group 2 (profile ids 39, 50, 65 and 84) made steady progress throughout the trial, as evidenced by increasing numbers of successful attempts without prompting. Figure 9 (middle) plots an exemplar from this group.

Children in Group 3 (profile ids 41 and 42) require prompting to learn across the trial, but even so, achieved increasing successful tasks attempts and decreasing prompting. Figure 9 (right) plots an exemplar from this group.

Of those children who did not evidence a clear trend in successful task attempts, all but one achieved more successful attempts than fails in any given skill.

Overall, we found that 40% of failed skills (18 out of 45) are eventually achieved without prompting, and 13% are achieved with some level of prompting. Moreover, the majority of skills (62% – 98 of 158) begun at some level of prompting were subsequently achieved without prompting. From the quantitative outcomes, we can confirm that TOBY met two important goals of any education application: it enabled *success* (by all but one child); and *learning* (Groups 1, 2 & 3).

The third educational goal targetted by TOBY is generalization of skills to daily life, which is a stepping stone to increased independence. To assess this aspect we turn to the results for NET. Recall that NET tasks make up 90% of the total tasks. A tasks could be attempted and completed more than once depending on how the child had progressed. NET tasks are also required to be revisited after a two week interval to test skill retention, and we found nearly all revisited NET tasks were successful. Of 2119 completed attempts (all participants), 1746 tasks, or 82.40%, were NET. Figure 10 (left) plots the popularity of NET tasks for this trial. From this we infer NET tasks were both accessible to the parents, and pitched appropriately with their partner iPad tasks.

Usage Patterns, Trial 3

With more participants, and a real-world scenario, Trial 3 afforded an opportunity to observe how TOBY was adopted into daily life. To do this we instrumented the app to log the timestamps and ids of all task attempts, together with the initial and subsequent state of the syllabus (e.g., if a task completion unlocks a new task). This allows us to see, e.g., if there exist temporal patterns for different kinds of tasks. In addition, all navigation interactions are logged, thus allowing us to see how much time is spent reading the help guide (and which page), looking at reports, or playing rewards.

The chart at the top of Figure 10 (right) plots total amount of time spent actively using the TOBY Playpad for 32 families who spent at least 1 hour using the app, and uploaded their data to our server. The bottom chart plots proportions of time spent in different parts of the app. Those who spent a reasonable amount of time in the app (on the left) performed a mix of tasks, such as receptive and expressive matching, and NET; whereas those who spent the least time in the app, on the right, spent their time on help and management tasks.

Time spent on NET varies across users. This reflected different modes of use for NET tasks, which were subsequently uncovered through questionnaire. Some parents reported doing NET tasks with the iPad present as a reminder of the task parameters, and so they could supply feedback immediately. Others preferred to digest the NET task on the iPad, but then put it away when performing the task, because their child found it difficult to concentrate with the iPad present.

Some participants used rewards heavily (ids 56, 31, 61), and recorded less progress in other iPad activities. This may be because the children had difficulty focusing on the iPad tasks. Most parents spent at least 20% of their time in admin sections of the app (i.e., Help, Manage, Reports, and Plan), some as much as 60%. There are relatively few users who spent above 70% of their time on course content (ids 30, 22, 54, 28, 34, 50), but these tended to be the users who had spent proportionally the most time doing high-level iPad tasks e.g., (receptive and expressive matching), and less time doing NET.

For certain users, we noticed the time between first view of a NET task and feedback was too short for the task to actually have been performed. This explained the large number of task completions during some sessions–sometimes as many as 50 in one session. During de-briefing we discovered this behaviour was due to their believing their child was already able to do the task. We subsequently added the ability to skip NET tasks in this situation.

CONCLUSION

We have traced the evolution of TOBY Playpad, an iPad app for early intervention therapy for children with autism. Initially the primary design challenge appeared to be how to

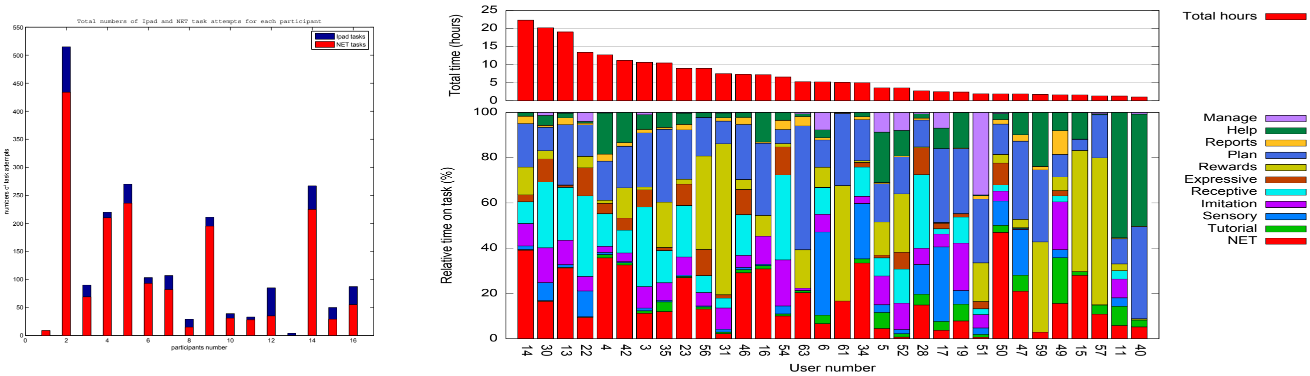

Figure 10. Left: Trial 2, Ipad and NET tasks. Right: Trial 3, total time using TOBY proportional breakdown, for 32 trial families.

encode therapist practice and deliver discrete trials. But as TOBY's scope expanded, and we observed and interacted with parents across a number of software trials, we came to realize that the major hurdle TOBY needed to jump was also its unique contribution: *teaching a parent how to teach*.

Unsurprisingly, this challenge led to the creation of an extensive, multimedia help resource. But it also caused the shift in TOBY's mixed-initiative daily lesson planning toward flexibility in task selection and pacing. This in turn led to the provision of richer progress reports by which to guide that parental intiative.

The parameters of three formal trials mirror the evolution of the software design, shifting, as they do, the onus toward parents, concluding with a large field trial that replicates the conditions of TOBY's ultimate deployment. The trials demonstrated clearly that TOBY engendered learning, but just as importantly, that the majority of parents came to understand the therapy process TOBY uses, and felt empowered because of their role in it. The immediate future will see TOBY's influence on learning outcomes tested in a large, independent clinical trial, extension of its syllabus to cover new skills and the 6-12 years range, and analysis of the feedback obtained from parents via various social media channels.

Acknowledgments

We would like to express our wholehearted appreciation to the children and parents participated in this study – without them, this research would have had never been possible. We wish them a bright future ahead. We also would like to acknowledge our speech therapist, Wendy Marshall and clinical psychologist, Darin Cains for their expertise and involvement in the development of TOBY and the trials.

REFERENCES

1. Bernard-Opitz, V., Sriram, N., and Nakhoda-Sapuan, S. Enhancing social problem solving in children with autism and normal children through computer-assisted instruction. *Journal of Autism and Developmental Disorders 31*, 4 (2001), 377–384.
2. Blocher, K., and Picard, R. Affective social quest: emotion recognition therapy for autistic children. In *In Socially Intelligent Agents-Creating Relationships with Computers and Robots*, Kluwer Academic Pub. (2002).
3. Bondy, A., Frost, L., and Bondy, A. *A picture's worth: PECS and other visual communication strategies in autism*. Woodbine House, 2001.
4. Bosseler, A., and Massaro, D. Development and evaluation of a computer-animated tutor for vocabulary and language learning in children with autism. *Jnl. of autism and developmental disorders 33*, 6 (2003), 653–672.
5. Hetzroni, O., and Tannous, J. Effects of a computer-based intervention program on the communicative functions of children with autism. *Journal of Autism and Developmental Disorders 34*, 2 (2004), 95–113.
6. Kientz, J., Hayes, G., Westeyn, T., Starner, T., and Abowd, G. Pervasive computing and autism: Assisting caregivers of children with special needs. *IEEE Pervasive Computing* (2007), 28–35.
7. Lorimer, P., Simpson, R., Smith Myles, B., and Ganz, J. The use of social stories as a preventative behavioral intervention in a home setting with a child with autism. *Jnl. of Positive Behavior Interventions 4*, 1 (2002), 53.
8. Lovaas, O. Behavioral treatment and normal educational and intellectual functioning in young autistic children. *Jnl. of consulting and clinical psych. 55*, 1 (1987), 3–9.
9. Schmidt, C., and Schmidt, M. Three-dimensional virtual learning environments for mediating social skills acquisition among individuals with autism spectrum disorders. In *Procs. of the Int. Conf. on Interaction Design and Children*, ACM (2008), 85–88.
10. Schreibman, L., Whalen, C., and Stahmer, A. The use of video priming to reduce disruptive transition behavior in children with autism. *Journal of Positive Behavior Interventions* (2000).
11. Stanton, C., Kahn Jr, P., Severson, R., Ruckert, J., and Gill, B. Robotic animals might aid in the social development of children with autism. In *Procs. of the ACM/IEEE Int. Conf. on Human Robot Interaction*, ACM (2008), 271–278.
12. Venkatesh, S., Greenhill, S., Phung, D., Adams, B., and Duong, T. Pervasive multimedia for autism intervention. *Pervasive and Mobile Computing* (2012).