printr: Exploring the Potential of Paper-based Tools in Low-resource Settings

Jay Chen
Computer Science
New York University
Abu Dhabi, UAE
jay.chen@cs.nyu.edu

Azza Abouzied
Computer Science
New York University
Abu Dhabi, UAE
azza@nyu.edu

David Hutchful
Grameen Foundation
Accra, Ghana
dhutchful@grameenfoundation.org

Joy Ming
Computer Science
Harvard University
Cambridge, MA, USA
jming@alumni.harvard.edu

Ishita Ghosh
School of Information
UC Berkeley
Berkeley, CA, USA
ighosh@ischool.berkeley.edu

ABSTRACT

Despite the recent push toward using information and communication technologies (ICTs) to replace paper-based workflows, there remain many barriers to designing appropriate and deployable ICT solutions that replace paper. As a result, paper tools such as forms, charts, and graphs continue to be widely used, especially in developing regions. While paper is not without its drawbacks, its advantages are especially relevant to low-resource settings as paper tools require only a fraction of the development, deployment, and operational costs of software apps. In this paper, we investigate how paper tools can be improved and combined with ICTs so that low-resource organizations working in developing regions can benefit from the advantages of both types of tools. We perform an exploration of existing tools to design, `printr`, a system that integrates into existing paper-based workflows by allowing an organization to rapidly generate paper tools that can perform some functions typically associated with computation — addition, subtraction, lookup, visual feedback, and visualization — without requiring the introduction of an ICT at the point of use. We compare two paper tools that `printr` produces with two mobile phone apps developed by a large NGO in Ghana and find comparable user performance between apps and generated paper tools.

CCS Concepts

•Human-centered computing → Empirical studies in HCI;

ICTD '16, June 03 - 06, 2016, Ann Arbor, MI, USA

ACM 978-1-4503-4306-0/16/06. . . $15.00
DOI: http://dx.doi.org/10.1145/2909609.2909649

Keywords

Information communication technology for development; ICTD; paper; paper-based tools; paper-based workflows

1. INTRODUCTION

Even as computing devices proliferate and we attempt to replace paper in professional and personal ecosystems with software applications, it is undeniable that paper continues to be pervasive. Paper is light, low-cost, familiar, accessible, easy to use, intuitive to manipulate, and convenient to distribute [27, 28, 18]. These compelling advantages in addition to the enduring legacy of paper, especially amongst paper-based workflows, provide strong inertia against switching to more technologically sophisticated replacements. In low-resource settings where infrastructure such as stable electricity, computer hardware and software, and skilled human capital are deficient, paper is often the preferred choice for data collection, data storage, and information dissemination.

Still, paper has its drawbacks. When compared to information and communication technology (ICT) alternatives, its content is less mutable once inscribed, paper lacks computational feedback, and data collected with it is difficult to digitize [5]. ICTs designed to replace paper characteristically exhibit an inverse profile of the weaknesses and *strengths* of paper because they are unable to capitalize on the accessibility and tangibility of paper [18]. The tensions that arise between the ubiquitousness of paper and the desirability of digitized data in paper-digital workflows highlight the challenges of combining paper and digital technologies [10]. In this work, we investigate the following research question:

Rather than a wholesale replacement of existing paper-based workflows with digital tools, can we augment paper in these workflows with ICT-like capabilities?

To answer this question, we consider the idea of using paper as the platform and implementing functionality typically associated with ICTs using paper instead. We define the term "paper tool" as broadly as possible to include plain, bound, or laminated paper with printed colored content and related stationary items such as pens or pencils, stickers, etc. Many existing paper tools demonstrate how paper by itself can provide limited forms of computational and visual feedback. Paper graphs, for instance, can provide simple computation in the form of lookups and a visual representation of relationships. Nomographs use scales to represent a three-variable equation that can be paired with a pen to, for example, determine ideal body mass [32]. Partographs can provide predictive feedback to midwives and they are cheap, effective, and popular in low-resource settings [13]. Paper passbooks in microfinance provide visual feedback of transactional information that derive from the affordances of paper itself [14].

These examples show that it is possible to design paper tools that provide some computational feedback, and they also show that such tools can be useful in low-resource settings where paper is cheaper, more familiar, and more accessible than its ICT counterparts. Our work transforms these tools into useful archetypes. We abstract the functions of several ICT systems used in low-resource settings into three broad categories of tasks: tracking, information lookup, and decision-making. We then construct several customizable paper tools for each category that accommodate different applications. We then developed `printr` to help *publishers* construct paper tools on a computer that are then used by the end user in the absence of ICTs. In this manner, the capabilities of the publisher and those of the end user match the technology profile of the artifacts that they are expected to use.

To construct a paper tool using `printr`, the publisher chooses a category that best represents the application's function, customizes one of the suggested paper tools, then prints and assembles copies of them for deployment. Our work contributes the following:

1. The exploration of customizable paper tools for three archetypal tasks: tracking, information lookup, and decision-making.
2. The design and implementation of, `printr`, a system for enabling low-resource organizations to rapidly generate paper tools that do not depend on ICTs at the point of use.
3. A comparative evaluation of `printr`'s paper tools against two existing smartphone applications that were designed and implemented by a large NGO in Ghana.

2. RELATED WORK

We begin our discussion of related technologies by roughly classifying ICTs along a spectrum from 'more like a computer' to 'more like paper'.

Paper-like devices include tablets and mobile devices such as PaperTab, PaperFold, and MorePhone [16, 15, 31]. These electronic devices look and feel like paper. Their key selling points are their paper-like properties, including portability, maneuverability, and physical flexibility (e.g. they can bend or fold), while still providing the full computational power of a tablet or phone. These devices however include some of the benefits of paper at a high monetary cost making them unsuitable for low-resource settings.

Paper computing tools attempt to use paper documents as interfaces to the digital world [19]. Paper computing includes systems that use paper to control technological environments, such as DigitalDesk, which projects graphical interfaces on a paper-based workspace, or PaperPoint, which allows users to control and annotate PowerPoint presentations using printed slide handouts [22, 36, 29]. Other paper computing tools include paper with printed visual markers such QR-codes or barcodes that link paper to electronic devices, specialized devices for digitizing hand-written notes and annotations[19, 1]. These technologies use paper in conjunction with a computer device as opposed to utilizing paper on its own.

Paper prototyping uses paper to model computers. Paper prototyping typically utilizes paper's accessibility and ability to be easily manipulated to simulate and test the usability of software UI designs in a "wizard of oz" experimental structure [4]. While computer devices are not integrated in the creation of the tool, the tool is paradoxically dependent on computer devices because it has no power when used alone.

Paper composites or combinations of paper with electronic materials, such as the paper and conductive thread combination in pulp-based computing, can give paper-based tools more computational power, but still exist to simulate computers using computer devices and electronic materials embedded in paper's fabric [6].

Paper tools are artifacts that are essentially constructed from paper. In the past, 'complex' paper tools have been embodied by nomographs that can help perform calculations. Very recently, paper has been revisited by efforts to produce extremely low cost or easy to produce tools. Researchers have found clever ways to augment the capabilities of paper through the application of biophysics to create a paper-based microscope [7] and even printed circuitry for programmable mechanisms [23]. We believe that systems like `printr` can similarly apply *computation* and *visualization* enhancements to paper.

2.1 Tools for Low-resource Settings

Looking specifically at tools developed for low-resource settings, most are heavily reliant on computer devices and fall in the *paper computing* category. In the research literature there are tools that recapture information on paper, such as CAM, which uses a camera to record and store data from paper logs in rural microfinance groups in India; ODK Scan, which applies computer vision techniques on small "snippets" of the content for a single form field; and Shreddr, a combination of computer vision, database, and crowdsourcing techniques to transmit and verify information from paper forms [25, 3, 9, 8, 5].

Other tools use technology to track the actual process for data input on paper-based forms such as Digital Slate, which uses the pen to paper writing motion to capture data in savings groups in India, and Partopen, which is a digital pen specifically developed for filling out the partograph in maternal labor [26, 33, 34]. Yet other tools use a mix of technology and paper in the information gathering process itself, such as Local Ground, a barcoded geospatial surveying tool [35]. These tools seek to combine paper with computing to leverage the benefits of both and demonstrate the recognition that some of paper's affordances are valuable enough to be worth preserving in ICT-based solutions.

Singh et al. examined the design of numeric paper forms that balance the tradeoff between ease-of-use among target NGO populations and machine readability to streamline data collection [30]. Guimbretière et al. explored the relationship between paper and digital documents [17] and also their affordances in the context of active learning in the classroom [21]. In addition to being tightly coupled with computer technology, these works are mainly concerned with information *collection*, including both retroactive and active data capture. In contrast, `printr` produces tools for data tracking, lookup, and decision making applications, which can provide computational and visual feedback to the user.

The closest related work to our evaluation is, e-IMCI, a PDA-based system designed by DeRenzi et al. for administering the Integrated Management of Childhood Illness (IMCI) decision making protocol [12]. While we do evaluate an IMCI decision making tool produced by `printr`, the focus of our work is to consider a new way of combining paper with ICTs in a manner that integrates into existing low-resource workflows rather than through a wholesale replacement of processes using ICTs.

3. printr

`printr` is an interactive system that is designed to allow low-resource organizations to rapidly generate customized paper tools.

To avoid confusion, we refer to the individual using printr to build paper tools as the *publisher* and the individual using the paper tools created by printr as as the *user* or *end-user*. Through a series of questions printr allows a publisher to specify tasks. At each stage of specification, printr presents the publisher with a set of potential paper tools that could be printed for use by the target end-user.

In developing printr, one major challenge was abstracting the ideas and structure of the paper tools we considered into generic primitives to generate new tools. We began our exploration of the design space by surveying online resources for existing health and finance tools of all forms. We then spent a total of six weeks in Ghana observing daily operations of a microfinance institution (MFI) in Tema, and observing medical practices and conducting interviews at three government district and children's hospitals in the greater Accra region. In total, we met with 10 local microfinance and health experts (professors, doctors, operations managers, nurses) as well as intermediaries (loan officers, community health volunteers), and NGOs (leadership, developers) to learn about existing tool use and practices.

We found that all stakeholders we interviewed generally viewed ICTs positively. The administrators and leadership of these organizations were interested in shifting toward mobile apps and ICT solutions for efficiency or modernization reasons while the loan officers and nurses liked mobile apps for productivity reasons (lighter than paper, easy to use). However, while the end users expressed the desire to modernize toward ICTs, they were still firmly attached to paper in practice. One particularly salient example that we encountered was the MFI's maintenance of all transactions in triplicate using a combination of paper ledgers, paper passbooks, and receipt slips. The MFI already had computers and a database and the loan officers already were using smartphones. When asked why they kept so many redundant paper-based systems the operation manager replied that the paper receipts gave their clients a sense of security and trust. Recent works have described in detail the perceptions and practices around paper-based MFI tools [14] and organization-level challenges associated with paper-digital workflows [10]. Our experiences largely reflect the these previous findings, but the goal of this paper is to explore an interesting design opportunity that leverages these observations.

After examining the tools that we discovered, we spent four weeks designing different paper-based primitives and considering the mappings between the type of task and the computational or visualization capabilities needed to complete the task. We spent four weeks designing different paper-based primitives and considering the mappings between the type of task and the computational or visualization capabilities needed to complete the task. We spent another two weeks considering the specification challenges for publishers to design tools quickly and easily. We assume that the capability profile of our publishers is that they are: 1) familiar with web applications, 2) able to follow on screen instructions, and 3) qualified to select the appropriate designs to fit their needs.

Eventually, we took a task-based approach and categorized tools into three different kinds of tasks: tracking, information lookup, and decision-making. We make no claims as to whether these are the "correct" categorizations or generalizable to all paper tools, but we found that these three categories captured a large subset of the tools that we encountered (both online and in our needs assessment) and those that we designed ourselves. We omit the multitude of paper tools we designed due to lack of space. Instead, we describe the first version of printr, the paper tools users can design using printr, and interesting computational primitives. Where it is possible, we relate our paper tools to existing analogous tools that we encountered in our needs assessment.

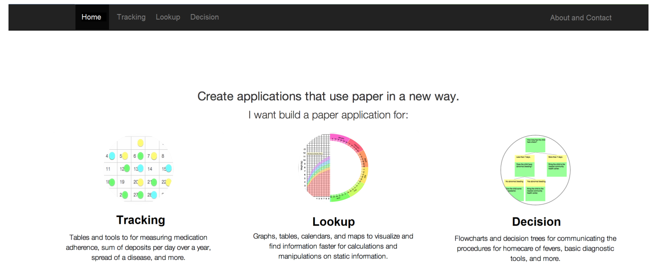

Figure 1: Screenshot of printr home page interface.

3.1 Implementation

printr is implemented as a web-based system with HTML / CSS and Javascript built utilizing the jQuery, Twitter Bootstrap, and D3 frameworks in 2134 lines of code. We developed printr using only client-side code so it can be used in settings with low connectivity. Figure 1 shows the printr prototype home screen where a publisher can select the type of tool they wish to design.

The specification process for each task category is different. As the publisher specifies their task, printr suggests multiple paper tools for the task if possible. printr accompanies each suggestion with information about the materials required and instructions for assembly and use. The publisher can then compare the possible tools (along with their associated materials and instructions) to select the one that best matches their specific design requirements.[1] For each of these types of tasks we describe how they are specified and the variations that printr currently produces.

3.2 Tracking

Tracking encompasses tasks that require periodic user input such marking the days a patient takes pills to visualize medication compliance, or tracking the days a microfinance savings customer makes the required deposit to calculate balance. Other possible use cases could include tracking daily caloric-intake over a month, periodicity of the menstrual cycle, daily journaling or mood tracking, etc.

In its most basic form, a single-column table captures most tracking tasks: each row represents a new entry. A time dimension can be introduced with a second time column that is pre-populated with the specified tracking frequency (days, weeks, months, etc). The tracking column itself can have different forms: for example, a checkbox form (or Yes / No) can be used to track whether medication was taken or not, an empty cell can be used to track generic numeric or string values. If the publisher wishes to track multiple events over time, printr checks that their time scales are compatible and then merges the separate multiple columns into a single table. Publishers can click a checkbox to disable automatic merging of tracking tasks if this is not desired. printr can also reorganize the table into a calendar form if tracking occurs daily for a month or more.

The publisher can also add additional columns that are pre-populated with data. For example, the publisher can add a running-sum column to a fixed-deposit savings tracking tool to provide information to the end-user on the total savings made.

Real-world use case: A Ghanaian microfinance institution (MFI) we encountered tracks deposits in paper passbooks (Figure 2a),

[1] In the future, we plan to have printr suggest only paper tools (and instruction-sets) for each task that match the publisher's specification of his/her end-users' capabilities (e.g. education-level, literacy, numeracy).

along with the dates of each deposit, amount deposited (even though they are fixed amounts) and the total savings made so far. Clients deposit a fixed amount each day for 31 days. The MFI Agents of the MFI insert this information and compute the total at each deposit. We replicated and produced two designs of these passbooks using `printr`.

Figure 2 shows the two tracking tool designs `printr` produces. The first tool has a column pre-populated with the running sum at each deposit. The second tool creates sticker values with the running sum at each deposit. Users track their deposits by placing the sticker with the smallest value on the calendar-date they made a deposit on. The tool, thus, visualizes when deposits are made and the savings made so far.

Specification of tracking tasks: To specify a tracking task, `printr` asks the publisher a series of questions to elicit the values of the input and the frequency and duration of the tracking. For each question, `printr` iterates on the publisher's answers and displays suggestions, allowing the publisher to visually see the effect of each additional constraint.

A screenshot of the `printr` Tracking interface is seen in Figure 3. If, for example, the user wants to track their medication consumption over a month, they would first be asked "What are you tracking?" They then are prompted to complete the sentence "I am tracking..." with "my medication consumption." At this step, the user will be presented with a table one column and only the task name, or "my medication consumption" as the first row. Then the user is asked to specify what types of input could be given at each point in the process. The types of inputs that the `printr` system currently supports includes tracking occurrence or a single numeric value each time. In the case of medication consumption, the user will specify that they are tracking "whether the event happens or not" and be presented with a tables that have "Yes / No" pre-populated for each row. Other outputs would include a two column tally of the number of times the event occurs and the number of times the event does not occur. And if the user had chosen a single numeric value for each time, the rows would be pre-populated with the cumulative sum for each row. Then, the user would be asked about the frequency and duration for which they will be tracking, allowing the tables to have an additional column that would be pre-populated with the time increments the user specifies. And the user would also be presented with a clock or calendar for tracking. In this case the user will mark the times in which the event occurs. Or the user could place a green or red sticker to better visualize regularity. From the suggested outputs, the publisher can either select and print a paper tracking tool, or continue to specify further details.

Some example outputs from `printr` medication tracking tasks are shown in Figure 4.

3.2.1 Variant: Planning

Low-income populations are typically engaged in the informal labor market, meaning their income inflows are irregular and variable, accumulated through a variety of financial instruments such as earnings from jobs and borrowing from sources of credit. `printr` can produce a tracking tool that helps end users check if their expected expenses exceed their expected incomes.

Figure 5 is a representation of this tool. It requires a base made from stiff paper such as cardboard. The base has two narrow slits through which paper strips can smoothly slide up or down. The horizontal axes indicate what each separate strip represents (income or expense) and the vertical axis indicates the amounts accumulated on each strip: an income of 20 cedis is represented by sketching-in two cells on the income stip.

The tool works as follows: For each day, the end-user sketches-in the income they expect to make on the income strip. For example, if the user expects to make 20 cedis[2] on day one, they sketch-in the cells on the income strip up to the 20 cedi mark on the vertical axis. After the user provides their expected income, they slide-up the income strips such that the sketched-in cells of a strip begin where the previous strip's sketched-in cells end (see Figure 5), i.e. the user stacks the sketched-in cells across the multiple income strips. The last income strip now marks the total expected income.

The user then plans out his expected expenses in a fashion similar to income planning but using the expense strips instead. The user can immediately detect at any point if expenses exceed income: the sketched-in cells of the expense strip will be above the sketched-in cells of the income strip.

After planning, the user can re-use the tool for income and expense tracking (erasing existing sketches or sketching-in unexpected incomes/expenses) and re-adjusting the strips to determine if they can make ends meet or not.

3.3 Information Lookup

Information lookup includes tasks such as calculating an individual's body mass index (BMI) or determining if the BMI falls within a healthy or unhealthy range, finding the proper drug dose for a given age and weight, etc.

Like tracking, the simplest information lookup paper tool consists of lookup-tables (Figure 6). The example in Figure 6a shows multi-valued inputs acting as keys to multi-valued outputs. Formula or calculation based lookups of discrete values can also be represented in a similar way using tables with pre-populated inputs and outputs for a given formula. Category lookups may be construed as tables where the lookup keys and/or lookup values are a set of categories.

When the number of lookup keys/values are small, simple *relational* lookup tables can be quite effective. However, for larger data sets, these tables can become too unwieldy or too complex. `printr` can re-organize lookup tables with two to three keys in a matrix instead of a table, this reduces the repetition of partial keys (Figure 6b). `printr` can also encode categorical values as colored cells to represent multi-valued values on a single cell.

Specification of lookup tasks: To specify a lookup task, the publisher provides the inputs and outputs or functions that calculate an output from any input.

3.3.1 Variant: Slide rules

In its most basic form, this information lookup paper tool consists of two (or more) fixed paper strips — a top input strip and bottom output strips — that sandwich a sliding paper strip with a *cursor* and one or more *pointers*. The user typically slides the central paper strip to align its cursor (a special pointer designated for input) with an input value on the top input strip. The remaining pointers point to different output values on the bottom strips. If the input values fall on a scale that 'wraps-around' (e.g. calendar days, time, integers modulo n), or the input values are categorical, `printr` re-organizes the linear slide rule onto a circular slide rule (a wheel) where a fixed, larger outer circle consists of both input and output values and a smaller inner circle consists of a cursor and output pointers.

Real-world use case: Tables and charts are used in Ghana to perform lookups on pre-populated information. A large NGO in Ghana created a smartphone app to help nurses quickly estimate pregnancy trimesters; Figure 8a is a snapshot of this smartphone app.

[2] A cedi is roughly 30 US cents.

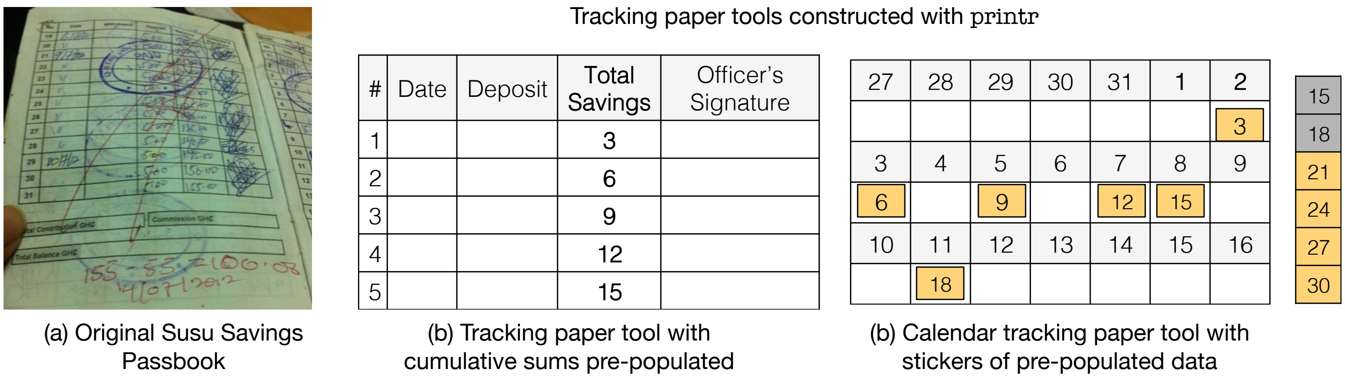

Figure 2: *Tracking example.* (a) The savings passbook of a Ghanian microfinance institution, (b) printr's improved passbook includes a column with pre-populated cumulative sums, (c) an alternative printr output uses a calendar layout and stickers.

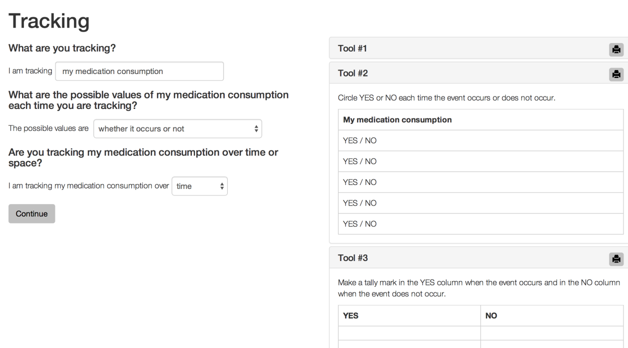

Figure 3: Screenshot of Tracking interface. On the left, the publisher specifies a tracking tool by stepping through a series of questions. Suggested tracking tools appear on the right.

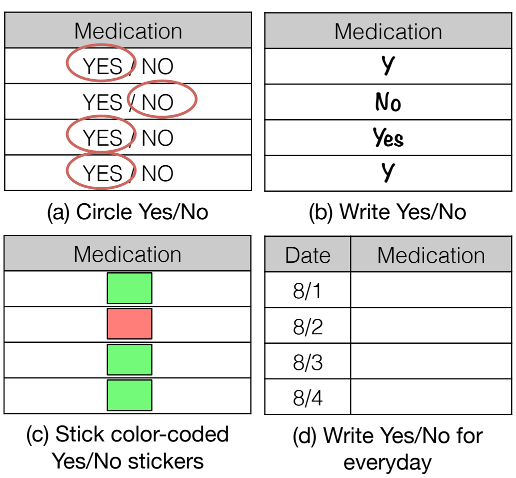

Figure 4: *Tracking example.* Example outputs produced with printr to track whether medication was taken or not on a daily basis over a month.

Specification of the trimester estimation wheel: Using printr, we constructed a paper version of the smartphone app by specifying the position of the cursor labeled 'Last Menses' to be at '0' and the position of three output pointers, 'First', 'Second' and 'Third' trimester to be at positions '90', '180' and '280' on a scale of 365 days. printr constructs the (circular) slide rule with the cursor and pointers correctly positioned. We specify both the input and output to be dates of the calendar year: printr merges these onto a single circular scale for the outer-circle). Figure 7 shows the user interface in printr for producing a circular slide rule. Figure 8b shows a circular slide rule that we produced with printr.

We evaluate the effectiveness of this paper tool against the NGO's smartphone app in the Evaluation Section.

3.3.2 *Variant: Viewfinders*

Large tables, complex charts, and colorful graphs can be intimidating or hard to use, especially for users have less experience with such tools. We make a paper viewfinder enhancement to paper graphs and tables that helps users focus on the relevant information.

Real-world use case: In Ghana, nurses have a standard government issued chart that can be used to determine whether a child is at a healthy weight-age balance based on the WHO regulations of child growth percentiles. This chart also tracks the child's weight over 24 months, accompanied by a legend to indicate the underweight or overweight regions (Figure 9a).

Figure 9b shows the printr produced viewfinder tool along with instructions on how to assemble the viewfinder tool (Figure 9c). This tool resembles a two-dimensional slide rule. The tool consists of a horizontal slider that users slide right or left to change the value on the x-axis. For the healthy weight-age graph, sliding the horizontal slider changes the age of the child. The graph itself is printed on a vertical slider that can be pulled up or down to change the value on the y-axis. Finally an outer fixed envelope encases and holds together the sliders: this envelope reduces the complexity of the underlying graph by hiding unnecessary information and focuses the attention on the user on whether a given weight/age lies in a healthy range or not. The envelope has windows to reveal the different x and y axis values as the user slides the horizontal or vertical sliders.

Specification of view finders: To construct a viewfinder paper tool with printr, the publisher provides a base graph as an image and marks the rectangular regions of the graph that contain the graph and the axis labels.

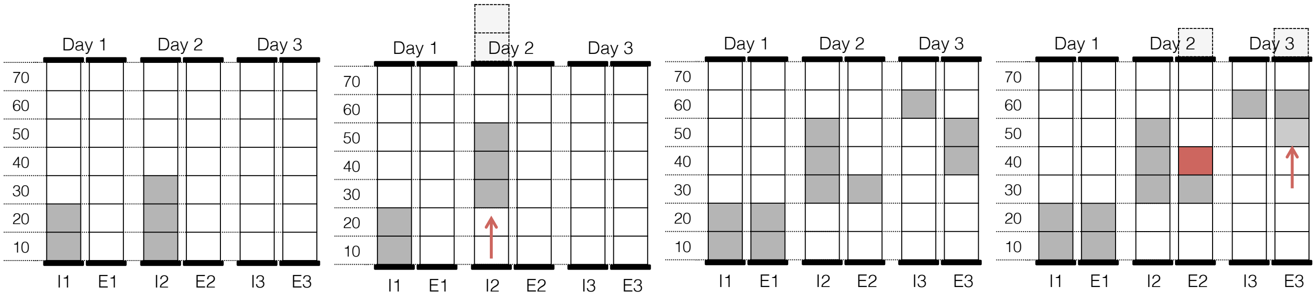

(a) Planning

(b) Slide to stack incomes

(c) Income meets expenses

(d) New expense requires re-adjustment

Figure 5: *Planning example.* In the planning phase, (a) the user sketches-in expected income and expenses for each day (b) the user adjusts strips to compare total expenses to total income (b), an example of a completed tool with both income and expenses populated and shifted (c). In the case of an additional expense, the user adjusts the strips to maintain the "stacked" invariant with the previous day's expense (d).

Inputs			Outputs	
Age	Weight	Height	BMI	Status
Adult	40	130	23.67	Normal
Adult	50	130	29.59	Over
Adult	60	130	35.50	Over
Adult	70	130	41.42	Obese

(a) Relational Layout

Height (cm) \ Weight (kg)	40	50	60	70
130	23.67	29.59	35.50	41.42
140	20.41	25.51	30.61	35.71
150	17.78	22.22	26.67	31.11
160	15.63	19.53	23.44	27.34

(b) Matrix Layout

Figure 6: *Lookup table example.* Different table-based representations of lookup tasks with different numbers of inputs and outputs.

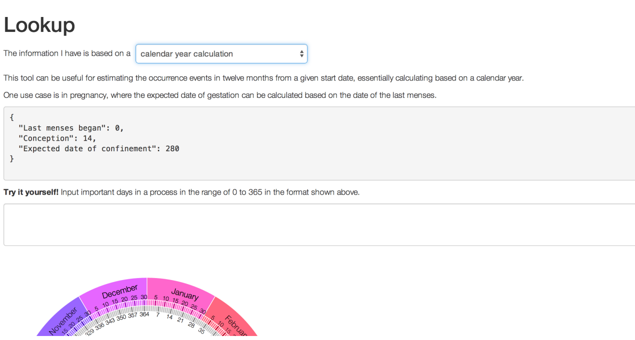

Figure 7: Screenshot of the Lookup interface when the user has selected "calendar year calculation" in the dropdown.

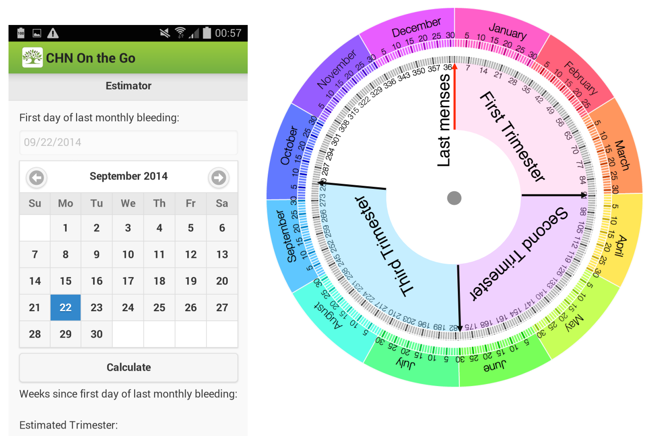

(a) Trimester Calculator App

(b) printr Trimester Circular Slide Rule

Figure 8: (a) An NGO's trimester calculator app. The nurse taps the left and right arrows to find the first day of last monthly bleeding on the calendar, taps the date on the calendar, and then taps the "Calculate" button to get the estimated trimester. (b) `printr`'s circular slide rule trimester lookup paper tool. The tool is used by turning the inner circle so that the red arrow lines up with the date of last menses. The trimester estimate is the region in the inner circle where the current date on the outer circle lines up.

3.4 Decision-making

Decision-making includes tasks that follow a sequence of steps in a flowchart or a decision-tree. Example decision-making tools include diagnostic tools that help volunteer medical workers or nurses diagnose and treat non-emergency patients, tools that help customers choose a savings schemes, etc. Decision-making tasks are best abstracted with trees, where non-leaf nodes contain questions, different edges are different answers to a node's questions and lead to further non-leaf nodes (i.e. more questions) or leaf-nodes that contain decisions or action-plans.

Specification of decision-making tasks: In `printr` the publisher builds the decision tree by adding nodes with questions and connecting them to predecessor nodes via labeled answer edges. Figure 10 shows a screenshot of the decision tree interface as it is being populated. If, for example, the publisher wants to create a tool that helps parents decide how to respond to their child's fever, the publisher first creates a node with the question of "How long has the child had a fever?" From this question, the publisher creates another node connected to the first node with the option "less than seven days" and the instruction to "bring the child to the nearest community health center." The publisher continues this until no more questions remain and all paths end in action-plans.

If the entire decision tree fits on a single page, `printr` properly lays out the tree for the user to print.

3.4.1 *Variant: Decision Booklets*

Certain decision trees may be very long or contain detailed instructions or information. Being confined to a single A4 sized paper in this case requires drastic simplifications to accommodate everything into single page diagram. To address this problem, `printr` produces decision booklets. This is helpful if space beyond a sin-

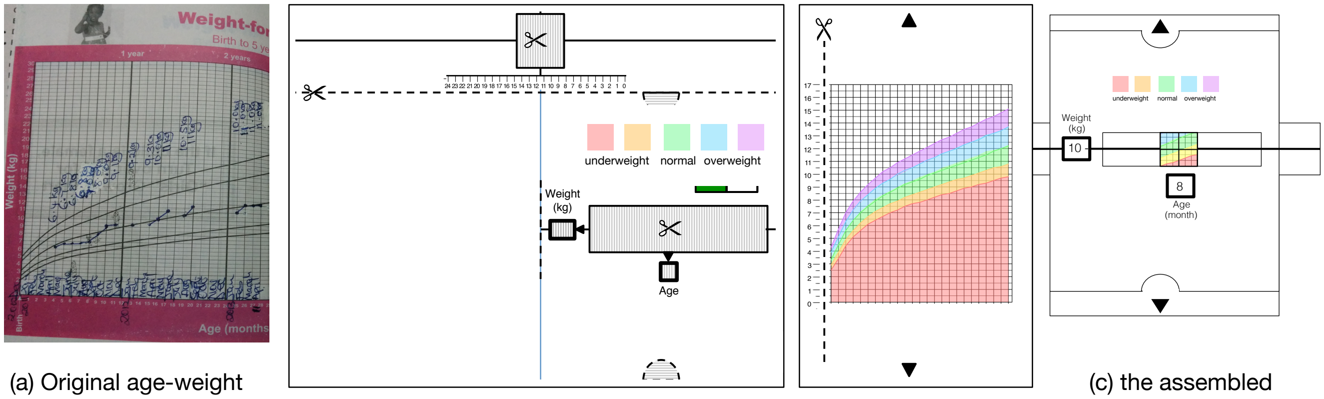

(a) Original age-weight WHO graph

(b) printr viewfinder tool with cut lines

(c) the assembled viewfinder tool

Figure 9: (a) The standard issue age-weight graph used in Ghana to determine a child's health, (b) `printr`'s viewfinder tool as printed, (c) `printr`'s viewfinder tool assembled.

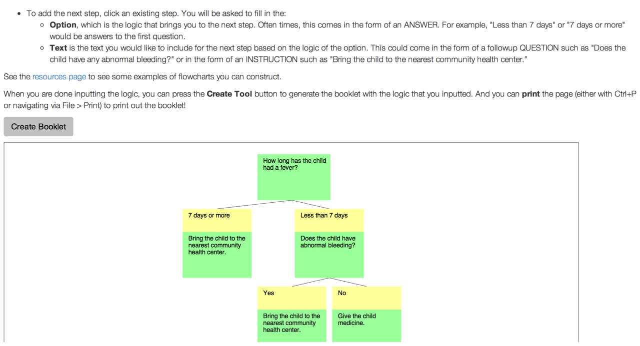

Figure 10: Screenshot of the Decision interface. The publisher specifies through the tree view and can convert this specification into a booklet using the "Create Booklet" button.

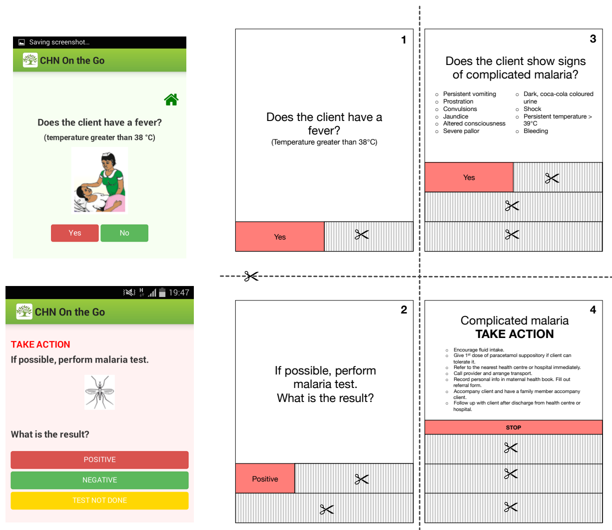

(a) IMCI Diagnosis app

(b) printr unassembled decision-making booklet

Figure 11: (a) IMCI Diagnosis app. The nurse follows the instructions on each screen and proceeds to the next screen by clicking on different buttons. (b) `printr`'s unassembled booklet. After cutting the pages as indicated by the cut lines, the cutouts are bound in the order of the top-left corner numbers. To make a decision (in this case, a step in the diagnosis), the nurse flips the tab of the option that corresponds to the decision and reads the content on the following page. The final instruction is a red tab labeled "STOP".

gle sheet of paper is needed for longer flowcharts, detailed descriptions, or large pictures.

Real-world use case: Nurses in Ghana use several decision trees, derived from the Integrated Management of Childhood Illnesses (IMCI) [2], for diagnosing common ailments in children and pregnant women. The same NGO that created the smartphone app for trimester estimation also developed a smartphone app to help nurses quickly and conveniently perform IMCI diagnoses. Figure 11a is a snapshot of the app.

Specification of the IMCI decision booklet: Using `printr`, we constructed a decision booklet with the same IMCI content as the NGO's smartphone app. Figure 11b) illustrates part of a decision booklet for IMCI diagnosis that `printr` produced. Once assembled, the user flips through the booklet using the appropriate tab visible under the diagnostic question at each decision point leading to the following step.

We evaluate the effectiveness of this paper tool against the NGO's smartphone app in the Evaluation Section.

4. EVALUATION

We conducted a preliminary user study comparing two paper tools (trimester estimation wheel and IMCI decision booklet) constructed using `printr` with the two existing smartphone applications for antenatal care we described in Section 3. The smartphone apps were designed by a large NGO in Ghana and ready for field trials with nurses. We partnered with this NGO to conduct our evaluation. While the smartphone apps designed by the NGO may not be "best-in-class" artifacts, they do represent real in-situ tools developed and deemed ready for deployment by a large NGO. Our goal was not to demonstrate "better" performance of `printr` tools, but rather to show that our tools perform comparably for the user at a fraction of the deployment cost/effort for the publisher. We observe task completion rates and times for each technology (paper vs. app) to compare the two types of interfaces.[3] We focus on evaluating the

[3]We do not use existing IMCI booklets as a baseline for comparison because they are, in principle, reproducible using `printr`.

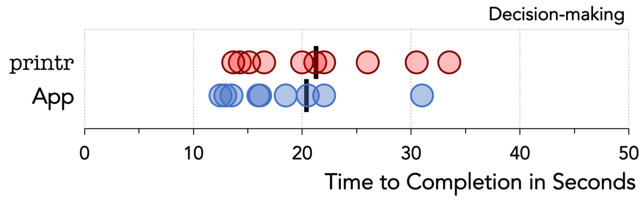

Figure 12: The average time to perform two decision-making tasks for each nurse on the smartphone app or with printr's paper tool.

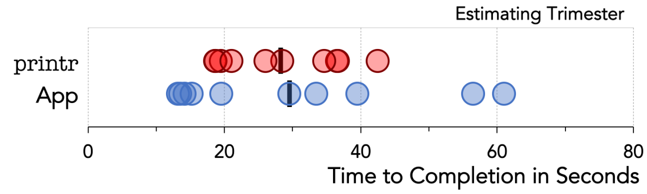

Figure 13: The average time to perform two information lookup tasks for each nurse on the smartphone app or with printr's paper tool.

tools produced by printr rather than printr itself since the printr interface is already a fairly easy to use web form for publishers (who are well educated and familiar with ICTs) even in its current prototypical state. Instead, we compare how well end users (who have less technical ability) perform using the paper tools produced by printr versus smartphone apps.

4.1 Participants and Methods

We recruited 10 nurses from the greater Accra region with the help of our NGO partner. The nurses were mostly female (8 out of 10) and their ages ranged from 25 to 31 years. The subjects were familiar with smartphone applications: all participants owned a smartphone. The participants have never used the NGO's antenatal-care application or printr's paper tools before.

We first conducted a 5 minute tutorial describing how to use the smartphone apps and the paper tools. We then asked subjects to complete two tasks for each tool or application. In each task, the experimenter plays the role of a client in need of antenatal care with pre-defined responses for each task.

Decision-making tasks: The diagnosis app or paper booklet guides the nurse through the examination and care of a client.

DM-Emergency. The nurse has to first determine the location of client (home or outreach clinic) from five different options (CHPS facility, health center or hospital). The nurse then determines whether the client is suffering from 'Oedema' or 'Heavy Bleeding' from five options (difficulty breathing, signs of shock or none). Each acute emergency leads to an action-plan.

DM-Fever. The nurse has to first determine whether the client has a fever (Yes/No question) and has either a negative malaria test or a positive malaria test with no complications from five options (positive test with complications and test not done). Each test result leads to an action-plan.

A timer starts when the nurse receives the smartphone (with the app loaded and ready) or paper tool and ends when the nurse reaches the action-plan. To reduce learning effects, a subject performs *two* DM-Emergency tasks with the smartphone app and *two* DM-Fever tasks with the paper tool or vice-versa. We randomized the presentation of tasks and tools across subjects. The time to task completion for each subject is recorded. After all four tasks were completed, we gave our participants a brief questionnaire comparing the smartphone app with the paper tool.

Information lookup tasks: The calculator app or paper tool helps the nurse estimate the trimester of a client using the date of her last menses. Again, to reduce learning effects, a subject estimates using the app or paper tool the trimester for *two* different dates. Like the previous task, a timer starts when the nurse receives the smartphone (with the app loaded and ready) or paper tool and ends when the nurse determines the trimester estimate. We randomized the presentation of tools across subjects. After all four tasks were completed, we gave our participants another brief questionnaire comparing the smartphone app with the paper tool.

4.2 Results

Figure 12 illustrates the average time for each nurse to perform decision-making tasks on the smartphone app versus tasks with printr's paper tool. Figure 13 illustrates the average time to perform information lookup tasks for each nurse on the smartphone app versus with printr's paper tool.

4.2.1 No significant difference in time of use

Our NGO partners believed that the smartphone app they built was more effective (w.r.t task completion times) than the paper tools. To test this hypothesis, we log-transformed task competition times to better approximate a normal distribution. We performed a repeated-measures ANOVA of completion times with tools as independent factors for both the decision-making and information-lookup tasks.

For the decision-making tasks, we found *no* significant main effect of tool used and we failed to reject the null hypothesis ($H_0 : \mu_{\texttt{printr}} = \mu_{\text{app}}$) at 95% ($F_{1,9} = 0.67, p = 0.43$). Similarly, for the information lookup tasks, we found *no* significant main effect of tool used and we failed to reject the null hypothesis at 95% ($F_{1,9} = 0.16, p = 0.69$).

Note that for both apps and paper tools, the total time to task completion includes time spent struggling with the tools. For example, with the decision-making app, two participants had a hard time finding the 'NO' button and one participant had difficulties scrolling down. We also observed participants struggling to find buttons in the calculator app: five participants could not easily find the 'calendar arrows' to change the month, or the 'calculate button'. We observed similar issues that slowed participants down with the paper tools. In particular most participants had difficulties aligning the arrows with the digits on the circles; a bigger circle with larger fonts can help solve such problems.

Though task accuracy and medical adherence were not the primary focus of our research, our performance results do complement DeRenzi et al.'s previous comparative evaluation between IMCI diagnosis when using PDAs (Personal Digital Assistants) and paper-based booklets [12].

4.2.2 Quantitative results

In our questionnaires we asked users to rate on 5-point Likert scale from strongly disagree (1) to strongly agree (5) (i) if they would use the app or paper tool again and (ii) if they found the app or paper tool easy to use. The mean rating for using the smartphone app again for decision making was 4.7: . This contrasts with a mean rating of 3.3 for printr's paper tool: . The mean rating, however, for using the smartphone again for information-lookup (estimating trimesters) was 3.6 () and the mean rating for printr's paper tool was 3.8 (). The mean rating for how easy users found the smartphone app for decision-making was also

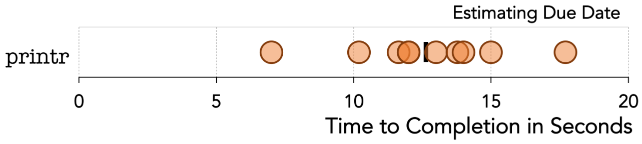

Figure 14: The average time to estimate due-date with `printr`'s paper tool.

high, 4.7 ( ), and contrasted with how easy they found the paper tool, 3.4 ( ). Users found the information-lookup app and lookup tool similar in terms of ease of use, 3.6 ( ) and 3 ( ).

4.2.3 Qualitative results

Users were given the opportunity to elaborate on their ratings. Most users (7/10) stated that they found the decision-making app "faster" than its paper-based counterpart — four of these users explicitly stated that they felt the paper tool was slower. This result is striking in the context of our quantitative results that show no significant effect of tool form on task completion time. Perhaps, users were exhibiting demand characteristics — they believed that the experimenters were hoping to prove the app was better[4] — or users found flipping pages to be somewhat laborious. One user commented that the paper tool is "not that easy because you need to be flipping," another commented that the paper tool "will be boring and tiring when I do it in a long time that's from one client to another." Users also found it difficult at times to properly line-up the arrows and calendar dates for the information-lookup paper tool: "very slow to wheel and look up the arrow up and down".

Despite the noted mechanical difficulties of using the paper tools, users recognized and identified their benefits and limitations. Several users remarked that paper tools can be used in the absence of smartphones, power or reliable internet connectivity. Users also pointed out that the paper tools can 'get burnt', lost, or easily torn. One user found smartphones phones more portable than the paper tools. Finally, many nurses appreciated the ability to access internet in general from a smartphone to perform other tasks and mentioned this in their feedback.

4.2.4 (Un)-Intentional benefits

Unlike ICTs, paper tools generally support computation by exposing their computation mechanics. In the `printr` Section, we described several computational primitives that paper tools support: for each of these primitives, the paper tool presents *more* information than necessary given the scope of the target application. For example, to support the tracking of cumulative savings, a running sum for each deposit is presented by the paper tool. While the primary purpose of the tool is to track savings and provide a total at each deposit, a side-effect is that users can figure out how many deposits are required to meet a certain savings goal. In the case of the trimester look-up tool, the nurses can determine the due date of a client by seeing where the third-trimester ends.

Without further training, we asked all 10 nurses to estimate the due-date of a client given her last menses date, and all 10 correctly estimated the due date in less than 20 seconds (see Figure 14). This calculation is not supported by the smartphone app without additional intentional effort, but follows naturally from the exposed computation mechanics of `printr` tools.[5]

[4]Demand characteristics have been demonstrated to be an especially dominant bias in developing contexts [11].

[5]This 'transparent' property has been described as one contributing factor to the persistence of paper by Ghosh et al. in prior work [14].

Table 1: Approximate cost (in USD) estimate of **one** "typical" smartphone application versus one "typical" `printr` tools. Training is typically performed in groups of 10 to 20 users.

Cost	**Smartphone app**	`printr` **tool**
Development & Testing	10-100 programmer hrs	10 publisher hrs
Production	$50(smartphone)	$1 (printing/assembly)
Training (local)	20 mins to 6 hrs	15 mins
Training (remote)	transportation	transportation
Deployment	transportation	transportation
Operation	battery charging	none
Updates	$1 (data plan)	production & deployment cost
Replacement	new phone every 3 yrs	new tool every 1 yr

5. DISCUSSION

The premise behind our work is that paper is cheaper and more accessible than ICTs. The implication being that if they could perform the same tasks, then in certain contexts paper tools would be substantially more cost-effective to deploy.

We perform a simple cost comparison between the smartphone applications our NGO partner developed (and intends to deploy) and the `printr` tools from our evaluation. We then discuss the limitations of paper, and identify situations where paper tools remain an interesting possibility for low-resource computational support at the point of use.

5.1 Cost Comparison

Table 1 estimates capital costs and operational costs for deployment of mobile apps and `printr` tools in various scenarios; it also lists any assumptions made in our cost estimation. Our NGO partner provided us with their development cost and time, and training time for the smartphone applications. We provided the publishing cost and time and training time for `printr` tools. These values are not meant to be comprehensive or perfectly accurate, but rather to give a sense of the comparative costs between mobile app development and `printr` tools.

These estimates are only for a single application or `printr` tool similar to the ones we evaluated. The deployment, maintenance, and update costs are drawn from local mobile, data subscription plans and transportation costs in Ghana. We also assume that end-users are already familiar with smartphones.

From Table 1 we see that the estimated development, production, and training costs for smartphone applications are at least an order of magnitude higher than `printr` tools. The developer of the smartphone app requires programming skills whereas the publisher only requires computer skills. Even if we assume that target users already own a compatible smartphone and can use a single phone for multiple applications, the development cost and time, and training time of a smartphone application still surpass those of a `printr` tool. The printing and assembly of paper-based tools is an expense, but at scale, the production, assembly, and deployment of paper tools could be automated as they currently are with existing paper tools.

Furthermore, while a smartphone does have dramatically more capabilities than paper, it also requires a power source for charging, and internet access for remote updates and data access [12].

5.2 Limitations of Paper

Despite the compelling cost and performance arguments for using a `printr` tool over an ICT solution, paper tools do have limitations that prevent them from being the best option in every situation. Paper has limited computational power, is limited by physical space, is limited by physical space and manipulation, and can be just as unintuitive as digital systems if designed poorly. Furthermore, paper tools require assembly (though this could be made part of the printing process at larger scales) and updates may require reprinting and distribution.

Somewhat outside of the properties of paper itself, but relevant to our context and applications is the enforcement of adherence of use (e.g. IMCI adherence [20]). While a digital solution can easily force users to adhere to a protocol through progress tracking and gated interfaces, paper tools do not have this obvious facility and can be prone to error. Another consideration worth bearing in mind is that ICTs are often perceived as aspirational in these contexts. As a result, ICT tools may be more enthusiastically greeted and therefore more consistently used [24].

Our goal with `printr` is to help low-resource organizations balance these limitations against the potential benefits of paper tools. If cost and training are considerations, paper tools may be more appropriate when the task: (1) requires simple and quick reading or writing from the user, (2) does not need much physical manipulation per step, (3) involves a limited amount of overall content, (4) can leverage pre-printed fields to reduce user effort, and (5) does not require a record of adherence.

6. CONCLUSION AND FUTURE WORK

In this paper we explored the potential of paper-based tools in addressing challenges in low-resource settings as an alternative to introducing a new mobile app. We find through assessments in microfinance and health institutions in Ghana that paper tools are still being used for many tasks despite the push toward ICTs. Through an exploration of existing and possible paper tool designs we developed a general taxonomy of paper tools and a system, `printr`, which can help low-resource organizations easily create tools that fit their target users. We show through a comparative evaluation with 10 nurses in Ghana that tools produced by `printr` offer comparable performance on diagnostic and calendar lookup tasks at only a fraction of the capital and operational costs. Moreover, the paper tools produced require little training and no ICT supporting infrastructure at the point of use.

There are many ways for `printr` to be extended, including improvements to specific modules or expansion of the set of primitives it provides. Future iterations could include more specialized tools such as maps for geospatial events in progress tracking or editable calendars for information lookup. Our goal is to add support for local languages and adjustment for different education levels. In the future, it would be worthwhile to explore whether the limitations of paper could be mitigated through techniques and determine the contexts in which paper is more suitable than an ICT counterpart.

7. REFERENCES

[1] Anoto: digital pen and paper technology. http://www.anoto.com.

[2] Integrated management of childhood illness. http://www.kznhealth.gov.za/chrp/documents/Guidelines/Guidelines%20National/IMCI/IMCI%20Chart%20Booklet.pdf.

[3] Y. Anokwa, C. Hartung, W. Brunette, G. Borriello, and A. Lerer. Open source data collection in the developing world. *Computer*, 42(10):97–99, 2009.

[4] C. D. Chandler, G. Lo, and A. K. Sinha. Multimodal theater: Extending low fidelity paper prototyping to multimodal applications. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '02, pages 874–875, New York, NY, USA, 2002. ACM.

[5] K. Chen, A. Kannan, Y. Yano, J. M. Hellerstein, and T. S. Parikh. Shreddr: Pipelined paper digitization for low-resource organizations. In *Proceedings of the 2Nd ACM Symposium on Computing for Development*, ACM DEV '12, pages 3:1–3:10, New York, NY, USA, 2012. ACM.

[6] M. Coelho, L. Hall, J. Berzowska, and P. Maes. Pulp-based computing: A framework for building computers out of paper. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '09, pages 3527–3528, New York, NY, USA, 2009. ACM.

[7] J. S. Cybulski, J. Clements, and M. Prakash. Foldscope: origami-based paper microscope. 2014.

[8] N. Dell, N. Breit, J. O. Wobbrock, and G. Borriello. Improving form-based data entry with image snippets. In *Proceedings of the 2013 Graphics Interface Conference*, pages 157–164. Canadian Information Processing Society, 2013.

[9] N. Dell, J. Crawford, N. Breit, T. Chaluco, A. Coelho, J. McCord, and G. Borriello. Integrating odk scan into the community health worker supply chain in mozambique. In *Proceedings of the Sixth International Conference on Information and Communication Technologies and Development: Full Papers-Volume 1*, pages 228–237. ACM, 2013.

[10] N. Dell, T. Perrier, N. Kumar, M. Lee, R. Powers, and G. Borriello. Paper-digital workflows in global development organizations. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pages 1659–1669. ACM, 2015.

[11] N. Dell, V. Vaidyanathan, I. Medhi, E. Cutrell, and W. Thies. Yours is better!: participant response bias in hci. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1321–1330. ACM, 2012.

[12] B. DeRenzi, N. Lesh, T. Parikh, C. Sims, W. Maokla, M. Chemba, Y. Hamisi, M. Mitchell, G. Borriello, et al. E-imci: Improving pediatric health care in low-income countries. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 753–762. ACM, 2008.

[13] A. Fawole, K. Hunyinbo, and D. Adekanle. Knowledge and utilization of the partograph among obstetric care givers in south west nigeria. *African journal of reproductive health*, 12(1):22–29, 2009.

[14] I. Ghosh, J. Chen, J. Ming, and A. Abouzied. The persistence of paper: a case study in microfinance from ghana. In *Proceedings of the Seventh International Conference on Information and Communication Technologies and Development*, page 13. ACM, 2015.

[15] A. Gomes, A. Nesbitt, and R. Vertegaal. Morephone: A study of actuated shape deformations for flexible thin-film smartphone notifications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 583–592, New York, NY, USA, 2013. ACM.

[16] A. Gomes and R. Vertegaal. Paperfold: A shape changing mobile device with multiple reconfigurable electrophoretic magnetic display tiles. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '14, pages 535–538, New York, NY, USA, 2014. ACM.

[17] F. Guimbretière. Paper augmented digital documents. In *Proceedings of the 16th annual ACM symposium on User interface software and technology*, pages 51–60. ACM, 2003.

[18] W. Johnson, H. Jellinek, L. Klotz, Jr., R. Rao, and S. K. Card. Bridging the paper and electronic worlds: The paper user interface. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems*, CHI '93, pages 507–512, New York, NY, USA, 1993. ACM.

[19] F. Kaplan and P. Jermann. Papercomp 2010: First international workshop on paper computing. In *Proceedings of the 12th ACM International Conference Adjunct Papers on Ubiquitous Computing - Adjunct*, Ubicomp '10 Adjunct, pages 507–510, New York, NY, USA, 2010. ACM.

[20] S. Lange, A. Mwisongo, and O. Mæstad. Why don't clinicians adhere more consistently to guidelines for the integrated management of childhood illness (imci)? *Social Science & Medicine*, 104:56–63, 2014.

[21] C. Liao, F. Guimbretière, R. Anderson, N. Linnell, C. Prince, and V. Razmov. Papercp: exploring the integration of physical and digital affordances for active learning. In *Human-Computer Interaction–INTERACT 2007*, pages 15–28. Springer, 2007.

[22] W. Newman and P. Wellner. A desk supporting computer-based interaction with paper documents. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '92, pages 587–592, New York, NY, USA, 1992. ACM.

[23] S. Olberding, S. Soto Ortega, K. Hildebrandt, and J. Steimle. Foldio: Digital fabrication of interactive and shape-changing objects with foldable printed electronics. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, pages 223–232. ACM, 2015.

[24] J. Pal, M. Lakshmanan, and K. Toyama. "my child will be respected": Parental perspectives on computers in rural india. In *Information and Communication Technologies and Development, 2007. ICTD 2007. International Conference on*, pages 1–9. IEEE, 2007.

[25] T. S. Parikh, P. Javid, S. K., K. Ghosh, and K. Toyama. Mobile phones and paper documents: Evaluating a new approach for capturing microfinance data in rural india. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pages 551–560, New York, NY, USA, 2006. ACM.

[26] A. L. Ratan, K. Toyama, S. Chakraborty, K. S. Ooi, M. Koenig, P. V. Chitnis, and M. Phiong. Managing microfinance with paper, pen and digital slate. In *Proceedings of the 4th ACM/IEEE International Conference on Information and Communication Technologies and Development*, ICTD '10, pages 37:1–37:11, New York, NY, USA, 2010. ACM.

[27] A. Sellen and R. Harper. Paper as an analytic resource for the design of new technologies. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, CHI '97, pages 319–326, New York, NY, USA, 1997. ACM.

[28] A. J. Sellen and R. H. Harper. *The Myth of the Paperless Office*. MIT Press, Cambridge, MA, USA, 2003.

[29] B. Signer and M. C. Norrie. Paperpoint: A paper-based presentation and interactive paper prototyping tool. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction*, TEI '07, pages 57–64, New York, NY, USA, 2007. ACM.

[30] G. Singh, L. Findlater, K. Toyama, S. Helmer, R. Gandhi, and R. Balakrishnan. Numeric paper forms for ngos. In *Proceedings of the 3rd International Conference on Information and Communication Technologies and Development*, ICTD'09, pages 406–416, Piscataway, NJ, USA, 2009. IEEE Press.

[31] A. Tarun, P. Wang, P. Strohmeier, A. Girouard, D. Reilly, and R. Vertegaal. Papertab: Tablets as thin and flexible as paper. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '13, pages 2881–2882, New York, NY, USA, 2013. ACM.

[32] A. Thomas, D. McKay, and M. Cutlip. A nomograph method for assessing body weight. *The American journal of clinical nutrition*, 29(3):302–304, 1976.

[33] H. Underwood. Partopen: enhancing the partograph with digital pen technology. In *CHI'12 Extended Abstracts on Human Factors in Computing Systems*, pages 1393–1398. ACM, 2012.

[34] H. Underwood, S. R. Sterling, and J. K. Bennett. The partopen in training and clinical use-two preliminary studies in kenya. In *HEALTHINF*, pages 112–121, 2013.

[35] S. Van Wart, K. J. Tsai, and T. Parikh. Local ground: a paper-based toolkit for documenting local geo-spatial knowledge. In *Proceedings of the First ACM Symposium on Computing for Development*, page 11. ACM, 2010.

[36] P. Wellner. Interacting with paper on the digitaldesk. *Commun. ACM*, 36(7):87–96, July 1993.