Child Growth Monitor - A game-changing app to detect malnutrition

Submitted by Jochen Moninger (Welthungerhilfe) on January 18, 2018 - 6:43am
Last revised by Web Producer on June 21, 2018 - 3:09pm.
Proposal Status: Awarded--Pending Funding

1. Executive Summary

The quick and accurate detection of malnutrition, especially with children under 5, saves lives!

The Child Growth Monitor is a fool-proof solution based on a mobile app using augmented reality in combination with AI/machine learning. The mobile app scans children under 5 in 3D point clouds and video data from children to extract anthropometric measurements like *height, head-circumference, middle-upper-arm-circumference, body volume / body composition / weight* to determine height and weight and therefore establishes the wasting rate. This is the main indicator for detecting moderate acute malnutrition (MAM) and severe acute malnutrition (SAM).

Measuring and following up malnourished children will be made easier, quicker, cheaper and more accurate.

- Mothers and governmental frontline workers often fail to detect severe malnutrition of children. As a result, they do not help the child in the right way.
- The magnitude of a nutrition crises both in emergencies and chronic hunger situations – is often blurred. This hinders a determined response by emergency workers as well as policy makers.

For aid agencies such as UNICEF or Welthungerhilfe, nutrition assessments will become much cheaper, quicker and more accurate. Immediate data and response may transform the way we deal with nutrition crisis.

For state services such as the Integrated Child Development Services in India, accuracy and accountability in managing nutrition data and responses is dramatically improved, so that no malnourished child is left behind.

- We make it OpenSource 8 GNU General Public License v3.09 to let everyone participate.
- https://github.com/Welthungerhilfe/ChildGrowthMonitor

2. Consortium Team

Welthungerhilfe specialises in fighting hunger and undernutrition. The CGM was conceptualised during an innovation camp in India in July of 2017. Five of the camp participants are now part of the passionate, interdisciplinary and efficient team. The team includes:

Joachim Schwarz: *Initiative Lead*, 20 years of experience in undernutrition, 5 years as Welthungerhilfe regional director in India

Markus Matiaschek: *Technical Lead, Software Solution and Coordination*, highly-motivated and with exceptional IT skills

Sweta Banerjee: *Nutrition Expert in India, Pilot Project Coordinator*, with extensive experience in managing acute undernutrition (CARE, Welthungerhilfe, UNICEF)

Christina Plassmann & Antje Blohm: *Branding, Marketing & Fundraising*, passionate members of the Welthungerhilfe Marketing Team

Jochen Moninger: *Innovation Manager* for Welthungerhilfe, expert on Africa

Markus Pohl: *Consultant for Digital Health*

As a scientific advisory board to the initiative we have already been able to win

- **Dr. Irmgard Jordan,** Justus Liebig University Giessen Institute of Nutritional Sciences International Nutrition
- **Miriam Silke Pietzsch,** Humanitarian & Development Professional - Public Health, Nutrition Security, Livelihoods, Operational Research

Collaboration has been established with the SMART team hosted in ACF Canada (Nutrition survey methodology), ACF Spain (photo diagnosis to prevent malnutrition) and Fighting Hunger Foundation FHF/ ACF India.

3. The project

We have developed a first version of the app. The app currently focuses on functions that are essential to collect data for our machine learning system. Those functions include: recognition of the child through a QR code, input for personal data, the scanning process, as well as an input for manual anthropometric measures.

In the coming months, we will collect up to 10,000 data sets of children through the app. The data sets will include 3D scans as well as the manually measured height, weight, middle upper arm- and head circumference of children. Those data sets will be used to "train" neural networks. Hence, data quality, of both the anthropometric data and the scans is of utmost importance to us at this stage.

Prototyping: Since interception we have produced a working prototype measuring body length through infrared and camera: The prototype recognizes 14 points on the body, i.e. shoulders ankles etc. and combines this with the 3D point cloud from the infrared scan. The app checks for updates of the neural network and downloads them automatically.

Technology: The necessary codes and data sources are published under the GNU General Public License v3.0 registered at Github. We have documented progress under: https://github.com/Welthungerhilfe/ChildGrowthMonitor

*Collected Data*set consist of gold standard anthropometric measurements of height, weight, MUAC and head-circumference from 12 healthy children, more than 700 3d scans of those children (~240 complete front/360°turn/back scans), more than 30.000 point clouds (min/avg/max per scan 1/30/918), each point cloud has a timestamp and up to 38.000 points, each point has XYZ float values in meters which amounts to 10 GB of zip compressed point cloud data.

Successful first round of field testing: Completed testing of Minimal Viable Product V.1

Funding: Since starting in August 2017 with the development of the app, we have been able to raise start-up funding of about 30,000 Euro as Innovation Award of Welthungerhilfe and 60,000 Euro from Deutsche Telecom CSR to develop the prototype and conduct the field testing.

4. Use Cases, User Stories and Activities

Some **160 million children worldwide** under five years of age are too small for their age. More than 50 million children are acutely malnourished, i.e. too light for their height. Half of these children live in Asia, a third in Africa.

The consequences of malnutrition are horrifying:

- 3 million children die from the consequences of malnutrition every year, i.e. one child every 12 seconds. **We can help to reduce this number.**
- Several hundred million people have suffered permanent cognitive and physical damage due to malnutrition in the critical 1,000 days (from conception to 2 years of age).

Measurements are time- and cost-intensive, being made with measuring sticks and scales that can be reliably operated only by experts. The availability of functioning tools is limited; data are sometimes fabricated, incorrectly transmitted and analysed only locally.

- Existing anthropometric measurement of malnutrition is complex and laborious. The method involves the use of scales to measure weight and the use of length boards to physically measure height/length of individual each child. This methodology requires highly-experienced staff, it is time-consuming and thus expensive. SMART surveys for example, which are best practice for such surveys take about one month. With the 3D-technology of the Child Growth monitor, it will be much easier and faster to determine a child's weight and length to be able to calculate his/her nutritional status.
- In emergency situations, such as the current crisis in South Sudan, Aid organizations often have difficulties to understand the magnitude of malnutrition, which is essential for a timely response.
- Nutrition Front-line workers, such as the Anganwadi worker in India, frequently fail to detect acute malnutrition due to wrong measurement or non-measurement of children (while faking the charts). This leads to prolonged malnutrition and death of children.

Solution

The Child Growth Monitor is a fool-proof solution based on a mobile app using augmented reality in combination with AI/machine learning. The mobile app scans children under 5 in 3D point clouds and video data from children to extract anthropometric measurements like *height, head-circumference, middle-upper-arm-circumference, body volume / body composition / weight* to determine height and weight and therefore establishes the wasting rate. This is the main indicator for detecting moderate acute malnutrition (MAM) and severe acute malnutrition (SAM).

- Measuring and following up malnourished children will be made easier, quicker, cheaper and more accurate.
- Mothers and governmental frontline workers often fail to detect severe malnutrition of children. As a result, they do not help the child in the right way.
- The magnitude of a nutrition crises both in emergencies and chronic hunger situations – is often blurred. This hinders a determined response by emergency workers as well as policy makers.

For aid agencies such as UNICEF or Welthungerhilfe, nutrition assessments will become much cheaper, quicker and more accurate. Immediate data and response may transform the way we deal with nutrition crisis.

For state services such as the Integrated Child Development Services in India, accuracy and accountability in managing nutrition data and responses is dramatically improved, so that no malnourished child is left behind.

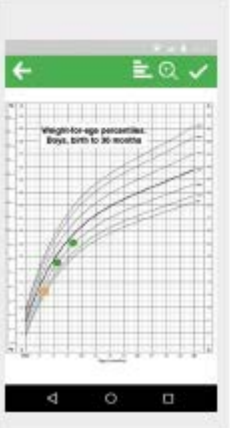

User Story

Mumbai slums are apocalyptic. We visited a mother of four, who had a seriously malnourished child that was treated in an ACF program. With her family of 6, she lives on 10 sqm, no window, no water, no sanitation facilities nearby, and directly bordering to a dark and narrow lane, which runs through the slums. The Child Growth Monitor is making the work in such areas much more efficient. With the Child Growth Monitor, the ACF team could identify and support more malnourished children and, at the same time, improve their analysis and ease their reporting.

Is this child malnourished or not? For her mother and even the nutrition expert it is difficult to tell! Measuring a child to determine malnutrition is not easy. Measurements are frequently incorrect.

Field Tests in Mumbai (8th March 2018) We are testing the App prototype and gather 3D data to train the artificial intelligence

Field testing in Mumbai, India, 6th-10th March 2018:

The purpose of the pre-test in Mumbai was to collect early feedback on the app from ACF field workers, and to assess the quality of the data we can expect both from the anthropometric measures as well as the scans. To test the app, we used a standardization test with 10 children under 5 years of age. Such tests are used in nutrition survey to ensure that all enumerators provide quality data. Five teams of ACF measured each of the 10 children twice - manually and through the app scan.

It was a critical exercise, in which we collected enough data to start building the architecture for the machine learning. The partnerships with ACF India, but also with ACF Canada, who helped in overseeing the standardization test, proofed as a great strategy.

Overall, the pre-test was a very successful exercise and the objectives were met

The importance of the concept of the Child Growth Monitor was re-confirmed by the ACF management and the field workers. It has the potential to make the work of the field workers so much easier, more accurate and more accountable.

The standardization test was successful and resulted in many quality data sets. Anthropometric data quality of the ACF team was very good, and we can be confident to get good data during the large field trial with 10,000 children. The involvement of Jana from ACF Canada was extremely helpful, i.e. to manage the process and to give important inputs on the CGM.

The scan process was good, but we realized that it needs to be further simplified, and special attention needs to be given to children that do not cooperate easily.

The relation with ACF was strengthened, better understanding of ACF programs and the needs of the field workers was created. All sides concluded to stay engaged and strengthen the partnership.

Many important feedbacks on the mobile app and the scanning process was received, mainly to simplify the scanning process and to improve user guidance. The scanning process of babies that cannot stand needs special attention.

5. Digital Health Technologies

The technological aspect is well documented and constantly updated on https://github.com/Welthungerhilfe/ChildGrowthMonitor .

We kindly refer you for details to this website.

The development objective is to create an open standard that can be tied to established platforms.

- Open source
- Software interface from smartmethodology.org

The project was started and is being administered as an open-source project. That means that the source code for the app, the back end and the machine-learning components are available under the open-source license GNU General Public License v3.0.

The data collected with the solution as well as the neuronal networks trained with the same will be secured in accordance with the need-to-know principle and current standards and methods.

Mobile App

The mobile app provides authenticated users an interface to scan children in 3D with consent of the parents and upload all collected data to the secure backend.

Currently the App works through person detection and pose estimation and overlaying the information of the position of 14 points on the body of the child with the 3D point cloud from the Tango API.

The next iteration for the Pilot will guide the user to scan the child in a way that a quick, accurate measurement can be taken. This will involve data of the camera pose, point clouds and RGB video.

Hardware requirements

Currently Google Project Tango devices only

In the future probably all devices with ARkit/ARcore capabilities (iPhone 6s and newer, 100 million Android devices)

Authentication

Users can authenticate themselves via username and password or Google OAuth. This enables access to download the latest neural network and upload data to Firebase Storage and Database.

User Flow

Also see this first UX Prototype

User eXperience (UX)

- Augmented Reality Userinterface guides through the scanning Process
- Scanning with instant visual feedback for the prototype the results of the scanning process will be shown only after input of the traditional measurements and only if more or less accurate

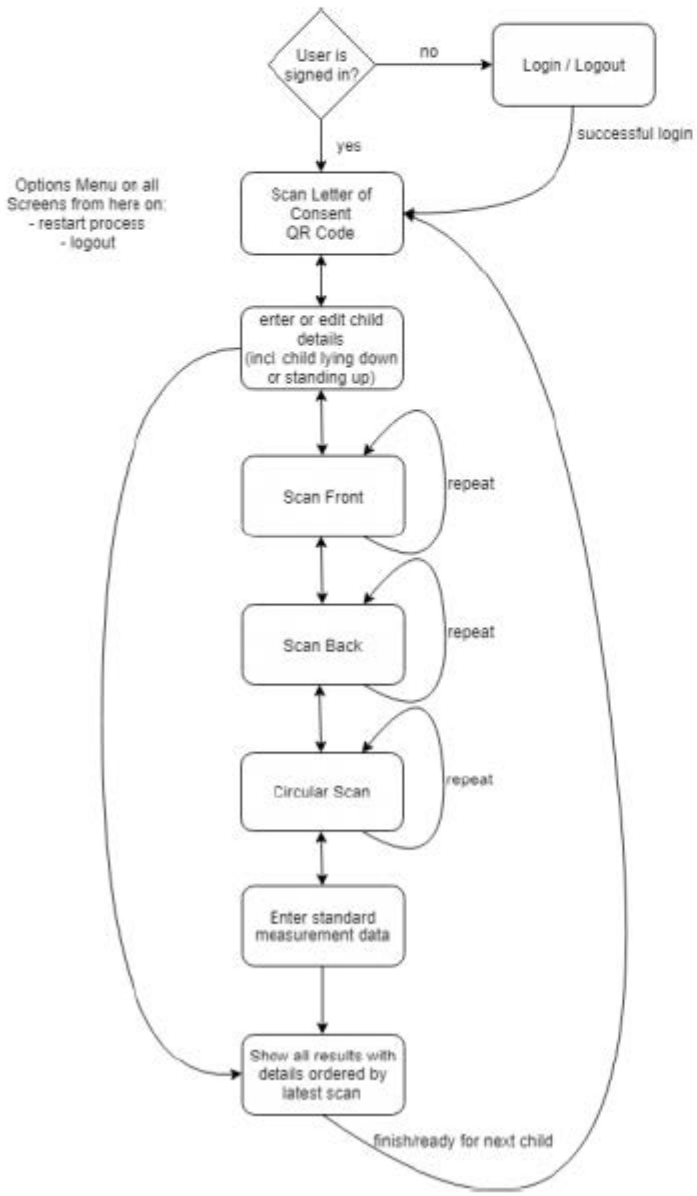

Backend

Backend is implemented in Google Firebase using Authentication, Database, Storage and Hosting for the Website.

Usermanagement

Users have to be activated by admin to download the current neural networks and upload data. Registration can be done via mobile app or the website via Firebase Functions.

Rights/Roles

Access to data is granted after scanning the key from a letter of consent of the parents.

Database

Firebase Database is used for structured data.

Storage

Storage is used for large objects such as rgb video and maybe point clouds.

Machine Learning v1.0

There are many possibilities for developing useful neural networks.

Predict height of a person

- An accurate prediction of the height of a human is priority number one. Goal is to do an 99,5% accurate prediction, so that we can measure a

---

child of 100cm height with an error of +/- 5mm.
- To reconstruct a 3d model of a child or of the skeleton is a non-trivial task using multiple point clouds of a moving child. Using a single point cloud isn't accurate enough.
- A promising approach is to input the point clouds, the device pose for camera position and rgb video into different neural networks, to do preprocessing or get the result.
- Helpful research has been done in
  - 3d point cloud segmentation through labeling the points
  - building a spatio-temporal graph for human pose detection

Predict weight of a person

- Predicting the weight of a person is the secondary goal to do "traditional" standardized measurements only using a smartphone without further hardware.

Classifying SAM, wasting, stunting, overweight

A promising approach could be to build a classifier to identify health issues. Downside to this is that without traditional measurements it is not possible to verify the decisions done by the classifier.

Current Machine Learning Solution

Human Pose Detection

Video Sequence -> Body Part Detection -> Spatio-Temporal Graph

-> Inference -> Pose Tracks

Person Detection

Bottom Up:

- fully convolutional Resnet-101 [He et al. 2016] -> Part Probability Scoremaps (heatmaps) aka. where are Shoulders, Knees, Wrists, Feet, …
- Discretize Scoremaps with non-maximum suppression (NMS)

TODO: Advantages of Top-Down approach?

Pose Estimation

The current resnet with 101 layers detects 12 joints + forehead and chin of a Person

Spatio- Temporal Graph G = (D, E)

Part Detections D

Spatial edges Es within one Frame

Temporal edges Et across frames

Provides Distance Features aka euclidean distance between body parts

--> Bottom Up Sparse Part Connectivity --> Euclidean Distance between Body Parts?

-

People Tracking

ResNet-101

6. Community Feedback

We work with the SMART teams of ACF Canada & Spain as well as the Fighting Hunger Foundation / ACF India on the data quality of the field trial, for getting feedback on the app, as well as integration with the global nutrition community.

The digital activists of the German Chaos Computer Club support us in further developing the technology.

We want to make the app part of India's health strategy and approach also governmental ICDS in India. With IT-development we plan that the CGM will be integrated into state-led programmes for managing acute undernutrition and not only NGO-led programmes.

In cooperation with partners we plan to organize datathons with partners to gain workforce and increase mutual learning.

To implement the existing requirements, we depend on additional specialists from the fields of augmented reality, 3D measurement and machine learning. Additional IT experts are already involved in technical implementation, especially in the fields of machine learning, 3D software, mobile development and user experience. Additional supporters are being recruited.

7. Self-assessment on the Global Good Maturity Model

Link to Self-Assessment: https://docs.google.com/spreadsheets/d/1vlTT5B1CyTXkt9g9cc17hC-US4rCEZJegeckx1_7KO8/edit?usp=sharing

Currently the app is in an early stage, therefore the self-assessment is currently quite low. However, once the field tests and data collection are completed, a higher maturity should be achieved quickly.

The Child Growth Monitor is globally scalable via several hundred governmental and non-governmental organisations. So far, all discussion partners have been enthusiastic, and many immediately identified it as a game changer in their work with undernourished children. With IT-development the CGM will be integrated into programmes for managing acute undernutrition (ACF India, governmental ICDS).

Once we have proven accuracy and functionality of the Child Growth Monitor app we will establish a scaling plan to have the solution tested and used in India and subsequently in African countries like South Sudan, Madagascar and Ethiopia.

8. Workplan, Project Deliverables & Schedule

Until 2018-08—with current financing

- Creating beta version; establishing scientific council (completed, February 2018)
- Pre-test of beta version in Mumbai with SMART and ACF India (completed, March 2018)
- Field test with 10,000 children in three Indian states utilising conventional methods (measuring by hand) and utilising our app in conjunction with ACF India (April to August 2018)
- Optimising algorithm; "make the case;" fundraising (August onwards)

From 2018-09 to 2020-03—with new financing

- Establishing an expanded team
- Continual IT development
- Field tests in two additional contexts (2019-02) in order to adapt CGM for children of other ethnicities
- Integrating the CGM in nutrition studies (beginning 2019-03)
- Preparing three pilot projects for managing acute undernutrition (by 2019-03)
- Pilot projects (beginning 2019-04) with approximately 100,000 children. Integrating CGM: (a) at ACF India; (b) with the state actor Integrated Child Development Services (ICDS) in the Indian Madhya Pradesh; (c) with the nutrition programme in refugee camps with regard to emergency aid
- Developing organisational structure and business model (by 2020-02)

Machine Learning Competition (Datathon)

Throughout the next 6 month we will collect data of

- ~10.000 children under 5 years of which around 40% will be MAM and 2-4% are SAM with gold standard anthropometric measurements with a refined scanning process that will provide between 6 to 30 point clouds per scan a confidence value with every XYZ point (XYZC) RGB or YUV video of the children > 250 GB of data
- Subsequent optimisation and adaptation of algorithm to determine for anthropometry: Height/Length, MUAC, Head-Circumference, Weight / Body Composition, SAM/MAM/Normal Classification from Video