Top Down Approach to Height Histogram Estimation of Biomass Sorghum in the Field

Jihui Jin, Gefen Kohavi, Zhi Ji and Avideh Zakhor; University of California, Berkeley; Berkeley, CA

Abstract

The rise of cheaper and more accurate genotyping techniques has lead to significant advances in understanding the genotype-phenotype map. However, this is currently bottlenecked by manually intensive or slow phenotype data collection. We propose an algorithm to automatically estimate the canopy height of a row of plants in field conditions in a single pass on a moving robot. A stereo sensor pointed down collects a series of stereo image pairs. The depth images are then converted to height-above-ground images to extract height contours. Separate height contours corresponding to each frame are then concatenated to construct a height contour representing one row of plants in the plot. Since the process is automated, data can be collected throughout the growing season with very little manual labor complementing the already abundantly available genotypic data. Using experimental data from seven plots, we show our proposed approach achieves a height estimation error of approximately 3.3%.

1. Introduction

The availability of cheaper and more accurate genotyping techniques in recent years has allowed for significant advances towards addressing growing challenges such as ensuring global food security and creating a sustainable biofuel source. Automatic extraction of biomass information such as crop height removes a bottleneck allowing geneticists to further understanding of the genotype-phenotype map. Current techniques for estimating plant height in the field are manually intensive and inaccurate. Thus, only a few measurements are taken within a relatively homogeneous plot a few times throughout the growing season.

High throughput imaging techniques have been proposed to automate the measuring process. Li *et. al.* [1] provide a survey of the various applications for plant phenotyping, such as machine vision, hyper-spectral remote sensing and three dimensional (3D) imaging. Paulus *et. al.* [2] discuss the advantages of 3D systems specifically comparing low cost systems such as stereo vision to more costly systems such as laser scanning. While laser scanners provide higher accuracy, they are expensive and also difficult to deploy in the fields.

Various algorithms have been developed to take advantage of 3D sensors for analyzing and measuring plant growth. Paproki *et. al.* [3][4] demonstrate an advanced segmentation and processing technique for 3D meshes of cotton plants. In order to create the high precision mesh, 64 images are taken at equally spaced intervals around the whole plant for multi-view reconstruction. Images are taken on a turn table and tripod in controlled conditions in the lab. Paulus *et. al.* [5] and Kjaer *et. al.* [6] present different approaches using high precision 3D laser scanners. While these works capture the plants with high fidelity, they are limited to laboratory controlled conditions and require manual data cap-


Figure 1: The camera coordinate frame is presented by the x, y, z axis. The robot coordinate frame is represented by the X, Y, Z axis. The robot is assumed to move in the X direction. Camera angle θ and height h are used to transform between the coordinates.

ture. Ngueyn *et. al.* [7] propose a structured light based approach for phenotypic measurement. They use five pairs of stereo cameras to 3D reconstruct smaller plants, with the eventual goal of automating for field deployment.

There have been a number of recent works focused on automating the phenotype data collection process in the field. This can be divided into two major categories: mobile and static measuring systems. In this paper, we focus on ground-based agricultural robots, even though quadcopter based systems could also be a viable approach. Agricultural robots are designed to collect data in field conditions. They face challenges such as noise, proximity of robot to plants, wind, weather and sunlight. They are typically equipped with sensors to measure and estimate plant growth metrics. Mueller-Sim *et. al.* [9] describe a ground-based agricultural robot capable of automatically collecting leaf area, leaf angle and stalk width in field settings. Jay *et. al.* [8] apply structure from motion to reconstruct 3D models of plants. Jay *et. al.* have demonstrated their algorithm on crops in field conditions using a specialized agricultural robot, extracting metrics such as plant height and leaf area. Similar to previous works, Jin and Zakhor [10] estimate stem width of crops from time-of-flight sensors mounted on an agricultural robot. Baharav *et. al.* [11] demonstrate a similar pipeline to estimate the width of sorghum using time-of-flight sensors and the height using stereo sensors.

While mobile agricultural robot based algorithms have been developed to collect individual plant metrics such as width and leaf angle, there has been less focus on height estimation. Existing methods of height estimation primarily use static measuring systems. For example, Phan *et. al.* [12] mount a 3D laser scanner at a high vantage point overlooking a field and estimate plant height from the 3D point cloud. Friedli *et. al.* [13] also extract canopy height using 3D laser scanners in fixed locations through-

Figure 2: Example depth image of dimension $H \times W$.

Figure 3: Point A is projected onto the image plane. z_A is stored for the depth image.

Figure 4: Two dimensional images $P_{xyz}^x(u,v)$, $P_{xyz}^y(u,v)$ and $P_{xyz}^z(u,v)$ are stacked together to form a length 3 vector of images.

out the field.

In this paper, we leverage the mobility and flexibility of agricultural robot systems to develop a height estimation algorithm around data collected from sensors mounted on robots capable of traversing between crop rows. The proximity to the canopy allows us to use cheaper commercial 3D stereo sensors over more expensive 3D laser scanners. We propose an automatic height estimation pipeline robust enough to handle crops from early to late stages of growth. The system is designed to have a dynamic range capable of measuring plants with a difference in height of up to 3 meters. Automatic collection of field metrics throughout the growing season offers geneticists a powerful tool to better understand the genotype-phenotype map.

The outline of this paper is as follows: we discuss the details of our pipeline in Section 2. Next, we explain the experimental setup for data collection designed to test the accuracy of our algorithms and analyze the results of our pipeline in Section 3. Finally, we provide some concluding remarks in Section 4.

2. Estimation Pipeline

The proposed pipeline requires a high moving vantage point and a commercial stereo camera as shown in Figure 1. The stereo camera is mounted on the mast of a moving robot pointed down. The robot moves in the space between two rows of plants and collects a series of images at a fixed rate as shown in Figure 6. The robot is assumed to move at a fixed speed and in one direction. The infrared images, depth images, mast height and camera angle are then passed into the pipeline to estimate the height of plants within each plot. The algorithm is designed to be modular and independent of the agricultural robot as long as the above conditions are met.

The overview of our proposed height estimation approach is as follows: the robot collects a series of depth images which are then converted into point clouds and filtered for only the points corresponding to the row directly in front of the robot. We reduce the point cloud constructed from the depth images to a one dimensional (1D) curves representing the canopy's height above ground. These height contours are then registered to construct a global height contour for one row of a plot. Individual plant heights are then extracted from the global height contour. In the following sections we describe each step in more detail.

2.1 Coordinate Transformation

We define the x, y and z axis as the camera reference coordinate system with the z axis pointing away from the camera and

the y axis pointing up as shown in Figure 1. We also define a robot coordinate frame as the $X, Y,$ and Z axis with the X axis aligned with the direction of the robot's movements along the row of plants, the Z axis perpendicular in the direction of the plants, and the Y axis pointing directly upward. Although the origin of the robot coordinate frame shifts in the real world between each new set of images, we can relate the Y and Z values of pixels between different sets of images as they correspond roughly to a plant's height above ground and distance from the robot respectively. Both coordinate frames are shown in Figure 1.

The stereo camera creates a depth image $I_d(u,v)$ of dimension $H \times W$ as seen in Figure 2. In order to predict the height of plants from the depth images, we must first convert it into a point cloud as shown in Figure 3. For pixel location (u_A, v_A) in the depth image, we determine the projected point A found at (x_A, y_A, z_A) in 3D space defined by the camera coordinate system. We set $I_d(u_A, v_A) = z_A$ and then repeat this for each pixel within the range $[1,H] \times [1,W]$. Then, using the depth image and camera intrinsics, we determine the x and y values corresponding to each pixel (u,v) to construct 2D images $I_x(u,v)$ and $I_y(u,v)$


(a) Y-Channel Image   (b) RGB Image

Figure 5: (a) $P_{XYZ}^Y(u,v)$, the Y-Channel of point cloud post transformation representing a height above ground image. The original depth image before transformation can be found in Figure 2; (b) The corresponding RGB image.

Figure 6: A top down view of the plot. We wish to only estimate heights on plants in the front row circled in green. Plants in further rows are filtered out.

in a similar fashion. Note that by definition $I_z(u,v) = I_d(u,v)$. We then concatenate the three images $I_x(u,v)$, $I_y(u,v)$ and $I_z(u,v)$ into $\overrightarrow{P_{xyz}}(u,v) = [I_x(u,v), I_y(u,v), I_z(u,v)]$, a vector of 2D images as seen in Figure 4. When referring to a specific image within the vector, we will denote the coordinate axis as a superscript, e.g. $P_{xyz}^x(u,v) = I_x(u,v)$. We also refer to the three images as the x, y, and z channels respectively.

To simplify the estimation of plant height, we convert point cloud $\overrightarrow{P_{xyz}}(u,v)$ in the xyz coordinate frame to $\overrightarrow{P_{XYZ}}(u,v)$ in the XYZ coordinate frame as seen in Figure 1. We use an unconventional coordinate frame for the camera's frame of reference to simplify the transformation between coordinates. We make the assumption that the robot moves primarily in one direction with little to no change in yaw. This reduces the transformation to a rotation about the x-axis by the camera angle θ, i.e. the angle between the z-axis and the Z-axis, followed by a translation in the Y axis by the mast height h, i.e. the distance between the origins of the two coordinate frames. For every pixel (u,v), $P_{xyz}^x(u,v) = P_{XYZ}^X(u,v)$, so we can reduce all coordinate manipulations to 2 dimensional transforms focusing only on the y,z to Y,Z transformations.

Let us rearrange the points $\overrightarrow{P_{xyz}}(u,v)$ and vectorize it as a $3 \times HW$ matrix P'_{yz} where the rows correspond to the y and z coordinate of every pixel respectively. We append a row of 1s so that the applied transformation is homogeneous. Likewise, if we vectorize $\overrightarrow{P_{XYZ}}(u,v)$, we get a second $3 \times HW$ matrix P'_{YZ}. We can relate P'_{yz} and P'_{YZ} with the following equation:

$$P'_{YZ} = \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & h \\ 0 & 0 & 1 \end{bmatrix} P'_{yz} \qquad (1)$$

The end result is a new point cloud $\overrightarrow{P_{XYZ}}(u,v) = [I_X(u,v), I_Y(u,v), I_Z(u,v)]$ with each pixel (u,v) now storing the points in 3D space with respect to the XYZ axis in Figure 1. The Y channel, or $P_{XYZ}^Y(u,v) = I_Y(u,v)$, corresponds to a "height-above-ground" image as seen in Figure 5a. We use this to calculate the canopy height from a single channel.

2.2 Height Contour

Once the images are in the correct coordinate frame, we extract the height contour of the canopy. In order to reduce estimation error, we focus on the front row only, or the row of plants


Figure 7: We approximate the location of the front row by its average Z value Z_{fr}. We use the radial distances highlighted in orange to determine which pixels correspond to the front row.


(a) (b)

(c)

Figure 8: An example of (a) the Y channel image $I_Y(u,v)$, (b) $I_r(u,v)$, and (c) $I_{filtered}^Y(u,v)$ the Y channel image after filtering for the front row only.


(a) (b)

(c)

Figure 9: The tallest pixel in each column highlighted in red on the (a) Y channel image and (b) RGB image respectively. (c) Height contour curve plotted over the physical plants

closest to the robot as seen in Figure 6. The row behind it suffers from error due to the degradation of sensor accuracy at further distances. Any plants beyond that typically belong to a separate plot which would corrupt the measurements of the plot of interest if included.

We detect the front row by estimating its distance Z_{fr} from the origin of the robot coordinate frame XYZ as shown in Figure 7. We approximate this distance by assuming that pixels of the plants in the front row are the closest to the camera as seen in Figure 7. We calculate the radial distance of every pixel using only the Y and Z channel of the point cloud $\overrightarrow{P_{XYZ}}(u,v)$. If we let h denote the mast height, then we can generate $I_r(u,v)$, an $H \times W$ radial image which stores the distance of each pixel (u,v) from the camera center:

$$I_r(u,v) = (P_{XYZ}^Z(u,v))^2 + (h - P_{XYZ}^Y(u,v))^2 \tag{2}$$

Each $I_r(u,v)$ now stores the Euclidean distance with respect to the Y and Z axis of the pixel to the camera origin as shown in Figure 8b. Then, for each column $v_k, k \in [1,W]$ in $I_r(u,v)$, we select a corresponding $u_k = \operatorname{argmin}_u I_r(u,v_k)$. This gives us the W pixel locations (u_k,v_k) closest to the camera corresponding to the front row. We then average the corresponding Z values of each pixel (u_k,v_k) as in Equation (3) to obtain Z_{fr}.

$$Z_{fr} = \frac{1}{W} \sum_{k=1}^{W} \overrightarrow{P_{XYZ}}(u_k, v_k, Z) \tag{3}$$

Any pixel u,v whose Z value $P_{XYZ}^Z(u,v)$ falls outside of a range centered at Z_{fr} are zeroed out in all three channels to result in 2D images $I_{filtered}^X(u,v)$, $I_{filtered}^Y(u,v)$, and $I_{filtered}^Z(u,v)$. An example of the Y channel after filtering out all points except those corresponding to the front row can be found in Figure 8c.

Since we are only interested in the height of the canopy, it is wasteful to track all points associated with the front row. Instead, we reduce the point cloud into a 1-D curve representing the top of the plants. Specifically, for every column $v_k, k \in [1,W]$ in $I_{filtered}^Y(u,v)$, we select $y_k = \max_u I_{filtered}^Y(u,v_k)$, the tallest pixel in the Y channel as seen in Figure 9a. This reduces an $H \times W \times 3$ point cloud to a length W vector of height measurements along the X direction. We refer to this vector as the height contour H. The physical representation of the height contour with respect to the plant row is shown in Figure 9 and an example graphical representation in Figure 10. In the next two subsections we describe two methods to extract plant height from height contours.

2.3 Height Estimation

The overall goal of height estimation is to accurately estimate histograms of plant height for each plot. Current genetic research maintains emphasis on comparison of plots with different gene characteristics. In particular, the histogram, mean, and variance of each plot is compared with each other to evaluate each genotype.

2.3.1 Method 1: Individual Frame Estimation

In this approach, we assume each plant results in a peak in the height contour, and use peak detection to find its height. We


Figure 10: The generated height contour (orange) with the smoothed height contour (blue). The local peaks detected from the algorithm are given as blue dots on the smoothed contour. The x-axis is in pixels and the y-axis is in meters.


Figure 11: The height contour with median filter. The red points are the local maximum that were remapped to the median filtered graph and represent the final predicted peaks. The x and y axes are the same as Figure 5.

then aggregate all of the heights obtained in this manner across all the height contours on a corresponding to depth images given row of plants to estimate the histogram of plant height for that row. Since each plot is seen in multiple depth frames, the overcounting of plants does not affect the shape of the histogram since nearly all plants are overcounted by a similar amount with the exception of plants near the edge of the row which appear in fewer frames than the rest.

For each height contour generated, we perform local maximum estimation as follows. Each plant is assumed to correlate with only one local maximum on the contour line. As a result, the goal is to find peaks on height contour H. To denoise H, a 1-dimensional 11-pixel median filtering is used to produce H_m. To reduce likelihood of detecting the same peak twice, the H_m is smoothed using a Hamming window, producing H_{ms}. The smoothing process causes the height of detected peaks to be slightly lower than expected as shown in Figure 10.

Local maximum detection occurs on the smoothed contour H_{ms}. For each x,y pair in the contour line, the algorithm decides if the point is a peak. A peak is defined as any single point larger that its nearest L neighbors where L depends on planting distances. For our setup, we choose $L = 20$. To prevent the ends of the contours from being falsely classified as peaks, an interval

Figure 12: The total height contour for a specific collection.

at the beginning and the end of the contour are ignored. This also ignores errors caused by the boundary effects of smoothing.

As mentioned earlier, smoothing underestimates the mean heights of detected peaks. To address this, we remap detected peak locations to values on the median filtered height contour H_m which does not suffer from the same artifacts. An example of this process on a height contour is shown in Figure 11.

The method described in this section over counts plants because each plant appears in multiple frames. With the exception of plants near the edge of a plot, all plants are recounted the same amount. With approximately 60 frames for each plot, under-counted plants represent a minority of all detected plants and hence do not significantly affect the histogram and average height estimated per plot.

2.3.2 Method 2: Registered Frame Estimation

The robot takes a series of overlapping images of a plot, with every plant appearing in more than one image. Our approach is to register and concatenate successive height contours H_{ms} into one global height contour for a given row in a plot, and to use the peaks in that as location of each plant. This alleviates the over-counting problem of method 1. The height contours do not vary significantly from frame to frame beyond a translation in the x direction and noise. To mitigate error in detecting shift caused by noise, we use the denoised and smoothed height contours H_{ms}. To determine x-shift, we compare two consecutive height contours and find the shift that minimizes the $\ell 1$ loss of the middle range of the contours.

$$X_s = \operatorname*{argmin}_{x} \| H_{ms}^i - H_{ms}^{i+1}{}_x \| \qquad (4)$$

where H_{ms}^i and H_{ms}^{i+1} are the i^{th} and $(i+1)^{st}$ height contours respectively and $H_{ms}^{i+1}{}_x$ represents the second height contour shifted by x pixels. This process is repeated for each subsequent height contour H_{ms} generated from each frame to construct a global height contour of the plot. Once the global contour of a plot is constructed, we can use a similar process to method 1 to estimate height. Because there is no overlap in a height contour for an entire plot, it is possible to count the number of detected peaks to estimate plant number.

3. Results

All data was collected at University of Illinois at Urbana-Champaign. An Intel RealSense R200 was mounted on a mast connected to a robot and moved along each plot at a fixed distance. For each frame RGB, infrared, and depth images are taken at a resolution of 640×480.

The sorghum is arranged in 3 meter by 3 meter plots. Each plot has a distinct genotype with mostly homogeneous plants within the plot. The sorghum is planted in four rows per plot with


Figure 13: Comparison of 2 height contours. The previous height contour (dotted orange line) is x-shifted by the algorithm to match the next height contour (blue line).


(a)

(b)

Figure 14: An example of the histograms (a) generated from method 1 and (b) with the ground truth. The difference scale in the y-axis is caused by the detection of the same plant across multiple images.

75 cm between each row. Sorghum plants are planted approximately every 6 cm resulting in approximately 50 plants per row and 200 plants per plot. A diagram of a plot is shown in Figure 6.

The robot traverses in the spacing between two rows and collects data from one row. The data is then processed by our pipeline to estimate a height histogram for each plot. Ground truth measurements were collected on 7 plots allowing us to compare the results of our algorithm.

Each plant in a plot was manually measured to create the ground truth. The list of measurements for a row of a plot creates a histogram of heights for a single collection. Table 1 shows the mean of histograms for 7 plots as estimated by method 1 and ground truth. As seen the error for all ranges in predicted and ground truth mean varies from 1 cm to 6 cm. Averaged over all plots, the mean error is 3.3%. The majority of the error in our algorithm is caused by an average overestimation of the predicted local peaks verses the ground truth peak heights.

Table 2 shows the performance of method 2 as compared to ground truth for both height estimation error as well as the number of plants. Averaged across all plots, the height estimation error for method 2 is 1.7% which is slightly lower than that of method 1 with 3.3%. The plant number estimate for method 2 has 18.9% error compared to ground truth. The underestimation of plant number is mostly caused by the smoothing process where two peaks on H_m could be detected as one on H_{ms}.

(a)

(b)

Figure 15: A comparison of the histogram from (a) method 2 with (b) the actual histogram of plants for a specific plot.

Collection	Method 1 Error %	Method 2 Error %
122437	4.0	2.0
122529	3.4	2.5
122511	2.0	2.6
122418	3.1	.04
122338	2.4	1.5
122455	5.3	2.9
122545	3.0	.18
Average	3.3	1.7

Table 3: Error comparison of methods 1 and 2.

Example histograms for methods 1 and 2 are shown in Figures 14 and 15 respectively. As seen, the estimated and ground truth histograms look similar to each other.

4. Conclusion

In this paper, we introduced two methods for estimating the histogram of a row of plants. Since the algorithms are run automatically, data collection can be done with little human intervention and can be done throughout the year and growth cycle. The algorithm can be used for different plants in addition to sorghum.

5. Acknowledgements

Collection	Method 1 Mean (m)	Ground Truth (m)	Difference in Mean (cm)	Error %
122437/478	1.11	1.07	4	3.3
122437/467	1.12	1.07	5	4.0
122529/519	1.09	1.05	4	3.4
122511/508	1.12	1.08	4	3.9
122511/501	1.09	1.08	1	0.5
122418/448	1.11	1.08	3	3.1
122338/418	1.08	1.10	-2	2.0
122338/407	1.14	1.10	4	3.6
122455/484	1.14	1.08	6	5.3
122545/535	1.03	1.07	-4	3.0
122545/538	1.12	1.07	5	5.1
Average				3.3

Table 1: Performance characterization of Method 1

Collection	Method 2 Plant count	Actual Plant count	Plant count error %	Method 2 Mean (m)	Ground Truth (m)	Method 2 Error %
122437	45	59	23.7	1.10	1.07	2.0
122529	35	49	28.6	1.07	1.05	2.5
122511	40	51	21.6	1.11	1.08	2.6
122418	42	47	10.6	1.08	1.08	.04
122338	44	47	6.4	1.12	1.10	1.5
122455	69	53	30.2	1.12	1.08	2.9
122545	49	55	10.9	1.06	1.07	.18
Average			18.9			1.7

Table 2: Performance characterization of Method 2, and comparison with method 1

References

[1] Lei Li, Qin Zhang, and Danfeng Huang, A review of imaging techniques for plant phenotyping, Sensors 14, no. 11, 20078-20111 (2014).

[2] Stefan Paulus, Jan Behmann, Anne-Katrin Mahlein, Lutz Plmer, and Heiner Kuhlmann, Low-cost 3D systems: suitable tools for plant phenotyping, Sensors 14, no. 2, 3001-3018 (2014).

[3] Anthony Paproki, Jurgen Fripp, Olivier Salvado, Xavier Sirault, Scott Berry, and Robert Furbank, Automated 3D segmentation and analysis of cotton plants, Digital Image Computing Techniques and Applications (DICTA), 2011 International Conference, IEEE, pg. 555-560. (2011).

[4] Anthony Paproki, Xavier Sirault, Scott Berry, Robert Furbank, and Jurgen Fripp, A novel mesh processing based technique for 3D plant analysis, BMC plant biology 12, no. 1, 63 (2012).

[5] Stefan Paulus, Henrik Schumann, Heiner Kuhlmann, and Jens Lon, High-precision laser scanning system for capturing 3D plant architecture and analysing growth of cereal plants, Biosystems Engineering 121, 1-11 (2014).

[6] Katrine Heinsvig Kjaer, and Carl-Otto Ottosen, 3D laser triangulation for plant phenotyping in challenging environments, Sensors 15, no. 6, 13533-13547 (2016).

[7] Thuy Tuong Nguyen, David C. Slaughter, Nelson Max, Julin N. Maloof, and Neelima Sinha, Structured light-based 3D reconstruction system for plants, Sensors 15, no. 8, 18587-18612 (2015).

[8] Sylvain Jay, Gilles Rabatel, Xavier Hadoux, Daniel Moura, and Nathalie Gorretta, In-field crop row phenotyping from 3D modeling performed using Structure from Motion, Computers and Electronics in Agriculture 110, pg. 70. (2015).

[9] Tim Mueller-Sim, Merritt Jenkins, Justin Abel, and George Kantor, The Robotanist: a ground-based agricultural robot for high-throughput crop phenotyping, Robotics and Automation (ICRA), 2017 IEEE International Conference, pg. 3634-3639. (2017).

[10] Jihui Jin, and Avideh Zakhor, Point Cloud Based Approach to Stem Width Extraction of Sorghum, Electronic Imaging 2017, no. 17. pg. 148-155. (2017).

[11] Tavor Baharav, Mohini Bariya, and Avideh Zakhor, In Situ Height and Width Estimation of Sorghum Plants from 2.5 d Infrared Images, Electronic Imaging 2017, no. 17, pg. 122-135. (2017).

[12] Anh Thu Thi Phan, Kazuyoshi Takahashi, Atsushi Rikimaru, and Yasuhiro Higuchi, Method for estimating rice plant height without ground surface detection using laser scanner measurement, Journal of Applied Remote Sensing 10, no. 4 046018-046018 (2016).

[13] Michael Friedli, Norbert Kirchgessner, Christoph Grieder, Frank Liebisch, Michael Mannale, and Achim Walter, Terrestrial 3D laser scanning to track the increase in canopy height of both monocot and dicot crop species under field conditions, Plant methods 12, no. 1, 9 (2016).

Author Biography

Jihui Jin received his BS degree in Electrical Engineering Computer Science at the University of California, Berkeley (2016). He is currently pursuing his PhD in Electrical Computer Engineering at Georgia Institute of Technology focusing on signal processing, machine learning and computer vision.

Gefen Kohavi is currently pursuing a degree in Computer Science at the University of California, Berkeley. He is currently working under Professor Avideh Zakhor using image processing and sensor data. He is also doing machine learning projects through Machine Learning at Berkeley.

Zhi Ji is currently an undergraduate student in Electrical Information Engineering at the University of Electronic Science and Technology of China(UESTC) as well as a research assistant in Media lab at UESTC. Previously, he was an exchange student in Electrical Engineering and Computer Science at University of California Berkeley and worked in video and image processing lab at Berkeley. His research interests are signal processing, image processing, machine learning, 3D sensors, and computer vision.

Avideh Zakhor is currently Qualcomm Chair and professor in EECS at U.C. Berkeley. Her areas of interest include theories and applications of signal, image and video processing and 3D computer vision. Prof. Zakhor received the B. S. degree from Caltech and the S. M. and Ph. D. degrees from MIT all in electrical engineering, in 1983, 1985, and 1987 respectively. She was a Hertz fellow from 1984 to 1988 and received the Presidential Young Investigators (PYI) award in 1992. In 2001, she was elected as IEEE fellow, and in 2018, she was elected as Electronic Imaging scientist of the year by SPIE.