



Portable Handheld Optical Window Inspection Device

This device allows field measurement of defects such as commercial aircraft windows.

John F. Kennedy Space Center, Florida

The Portable Handheld Optical Window Inspection Device (PHOWID) is a measurement system for imaging small defects (scratches, pits, micrometeor impacts, and the like) in the field. Designed primarily for window inspection, PHOWID attaches to a smooth surface with suction cups, and raster scans a small area with an optical pen in order to provide a three-dimensional image of the defect. PHOWID consists of a graphical user interface, motor control subsystem, scanning head, and interface electronics, as well as an integrated camera and user display that allows a user to locate minute defects before scanning. Noise levels are on the order of 60 μm . (1.5 μm).

PHOWID allows field measurement of defects that are usually done in the lab. It is small, light, and attaches directly to the test article in any orientation up to vertical. An operator can scan a defect and get useful engineering data in a matter of minutes. There is no need to make a mold impression for later lab analysis.

The PHOWID system components consist of a scanning head, motor control board, graphical user interface, and supporting electronics, and weighs approximately 5 pounds (2.3 kg). The scanning head consists of an x - y positioner, optical distance sensor, carrier plate, and vacuum supplied suction cups. The head has an integrated camera and LCD (liquid-crystal-display) screen that allows a user to easily position it above the defect. Vacuum is supplied to suction cups through a dual circuit system, which includes check valves to hold the head against the window or other smooth surface up to 90° inclination in the event that the vacuum source is lost. The graphical user interface (GUI) displays the surface image and allows the user to describe an area of interest to be scanned at various speeds and resolutions. The GUI displays the scanning progress in real time. Once a scan is completed, software provides the automated measurement capability to determine defect length, width, and depth and will store this information in a file.

The GUI also communicates with the motion control electronics. These electronics control the x - y positioning motors that move the optical sensor in the sensor head. These electronics also combine the depth and position data in real time and stream the data to the GUI. The associated electronics box contains the vacuum pumps, optical-sensor conditioning electronics, and power supplies.

Functionally, the user places PHOWID over the defect using the integrated camera. The user selects the desired area, resolution, and scan speed at the laptop GUI and initiates scanning. PHOWID has a depth range of 0.01 in. (0.25 mm). It has a noise floor better than 60 μm . (1.5 μm). Usable scan area is on the order of an inch square. Smallest resolution of the scan in the x - y direction is on the order of 300 μm . (7.6 μm).

This work was done by Curtis Ihlefeld and Adam Dokos of Kennedy Space Center and Bradley Burns of ASRC Aerospace Corporation. Further information is contained in a TSP (see page 1). KSC-13218

Salience Assignment for Multiple-Instance Data and Its Application to Crop Yield Prediction

Automated mapping of crops saves on survey time and improves map accuracy.

NASA's Jet Propulsion Laboratory, Pasadena, California

An algorithm was developed to generate crop yield predictions from orbital remote sensing observations, by analyzing thousands of pixels per county and the associated historical crop yield data for those counties. The algorithm determines which pixels contain which crop. Since each known yield value is associated with thousands of individual pixels, this is a “multiple instance” learning problem.

Because individual crop growth is related to the resulting yield, this relationship has been leveraged to identify pixels that are individually related to corn, wheat, cotton, and soybean yield. Those

that have the strongest relationship to a given crop's yield values are most likely to contain fields with that crop. Remote sensing time series data (a new observation every 8 days) was examined for each pixel, which contains information for that pixel's growth curve, peak greenness, and other relevant features.

An alternating-projection (AP) technique was used to first estimate the “salience” of each pixel, with respect to the given target (crop yield), and then those estimates were used to build a regression model that relates input data (remote sensing observations) to the target. This is achieved by constructing

an exemplar for each crop in each county that is a weighted average of all the pixels within the county; the pixels are weighted according to the salience values. The new regression model estimate then informs the next estimate of the salience values. By iterating between these two steps, the algorithm converges to a stable estimate of both the salience of each pixel and the regression model. The salience values indicate which pixels are most relevant to each crop under consideration.

This approach produces better estimates than an existing “primary instance” (PI) approach does. The PI ap-

proach assumes that each county contains a single canonical pixel for each crop (corn, cotton, soybean, etc.) and that the rest of the pixels in that county are noisy observations of the true one.

This work could ultimately provide automated mapping of crops that are

being grown, which could save agencies such as the U.S. Department of Agriculture a significant amount of money that is currently devoted to surveying fields to produce summaries of how much of each crop is being grown. Reliable early estimates of the likely volume

of production can significantly affect crop prices throughout the season.

This work was done by Kiri L. Wagstaff of Caltech and Terran Lane of the University of New Mexico for NASA's Jet Propulsion Laboratory. For more information, contact iaoffice@jpl.nasa.gov. NPO-45177

Speech Acquisition and Automatic Speech Recognition for Integrated Spacesuit Audio Systems

This interface also has applications in mobile phones, in-car devices, and home electronics and appliances.

John H. Glenn Research Center, Cleveland, Ohio

A voice-command human-machine interface system has been developed for spacesuit extravehicular activity (EVA) missions. A multichannel acoustic signal processing method has been created for distant speech acquisition in noisy and reverberant environments. This technology reduces noise by exploiting differences in the statistical nature of signal (i.e., speech) and noise that exists in the spatial and temporal domains. As a result, the automatic speech recognition (ASR) accuracy can be improved to the level at which crewmembers would find the speech interface useful. The developed speech human/machine interface will enable both crewmember usability and operational efficiency. It can enjoy a fast rate of data/text entry, small overall size, and can be lightweight. In addition,

this design will free the hands and eyes of a suited crewmember.

The system components and steps include beam forming/multi-channel noise reduction, single-channel noise reduction, speech feature extraction, feature transformation and normalization, feature compression, model adaptation, ASR HMM (Hidden Markov Model) training, and ASR decoding. A state-of-the-art phoneme recognizer can obtain an accuracy rate of 65 percent when the training and testing data are free of noise. When it is used in spacesuits, the rate drops to about 33 percent. With the developed microphone array speech-processing technologies, the performance is improved and the phoneme recognition accuracy rate rises to 44 percent. The recognizer can be further improved by combining

the microphone array and HMM model adaptation techniques and using speech samples collected from inside spacesuits. In addition, arithmetic complexity models for the major HMM-based ASR components were developed. They can help real-time ASR system designers select proper tasks when in the face of constraints in computational resources.

This work was done by Yiteng (Arden) Huang, Jingdong Chen, and Shaoyan (Sharyl) Chen of WEVOICE, Inc. for Glenn Research Center. Further information is contained in a TSP (see page 1).

Inquiries concerning rights for the commercial use of this invention should be addressed to NASA Glenn Research Center, Innovative Partnerships Office, Attn: Steve Fedor, Mail Stop 4-8, 21000 Brookpark Road, Cleveland, Ohio 44135. Refer to LEW-18534-1.

Predicting Long-Range Traversability From Short-Range Stereo-Derived Geometry

Learning-based software improves obstacle avoidance by robotic ground vehicles.

NASA's Jet Propulsion Laboratory, Pasadena, California

Based only on its appearance in imagery, this program uses close-range 3D terrain analysis to produce training data sufficient to estimate the traversability of terrain beyond 3D sensing range. This approach is called learning from stereo (LFS). In effect, the software transfers knowledge from middle distances, where 3D geometry provides training cues, into the far field where only appearance is available. This is a viable approach because the same obstacle classes, and sometimes the same obstacles, are typi-

cally present in the mid-field and the far-field. Learning thus extends the effective look-ahead distance of the sensors.

The baseline navigation software architecture in both the LAGR (Learning Applied to Ground Robotics) and MTP (Mars Technology Program) programs operates so that stereo image pairs are processed into range imagery, which is then converted to local elevation maps on a ground plane grid with cells roughly 20-cm square covering 5 to 10 m in front of the vehicle, depending on

camera height and resolution. The image and the map are the two basic coordinate systems used, but only pixels with nonzero stereo disparity can be placed into the map. Geometry-based traversability analysis heuristics are used to produce local, grid-based, 'traversability-cost' maps over the local map area, with a real number representing traversability in each map cell. The local elevation and cost maps are accumulated in a global map as the robot drives. Path planning algorithms for local obstacle