Machine Vision for Airport Runway Identification

Matthew Schubert[abc], Andrew J. Moore[a]*, Chester Dolph[ad], Glenn Woodell[ae]

[a]Computational Vision Lab, NASA Langley Research Center, Hampton, VA, USA
[b]National Institute of Aerospace, Hampton, VA, USA
[c]Christopher Newport University, Newport News, VA, USA
[d]Vision Lab, Old Dominion University, Norfolk, VA, USA
[e]Science Systems and Applications, Inc., Hampton, VA, USA

Acknowledgements

- NASA Atmospheric and Environmental Safety Technologies program
- NASA Safe Autonomous Systems Operations program
- Rockwell Collins

Significance

Incorrect runway problem

- Over 500 incidents since 1999 (FAA)
- No injuries or fatalities, but close calls

Plane lands at wrong West Virginia Airport; No one injured

Boeing 747 lands at wrong airport in Kansas

Geolocation technologies

- Radio (e.g. VOR)
- GPS
- Gyros (inertial)
- Human vision

Source: FAA

Possible: Warning of Incorrect Landing

Night Approaches at 3 Airports – Very Strong Signal

Source: NASA

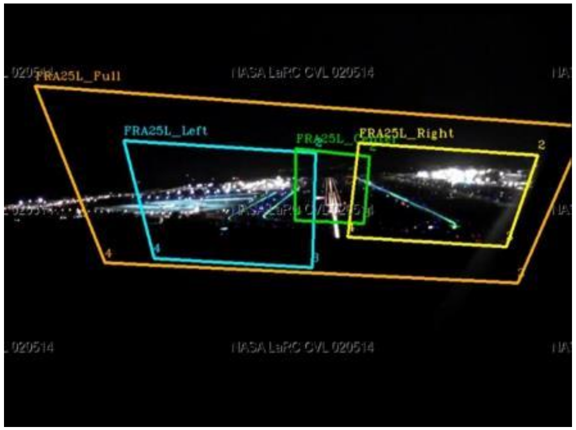

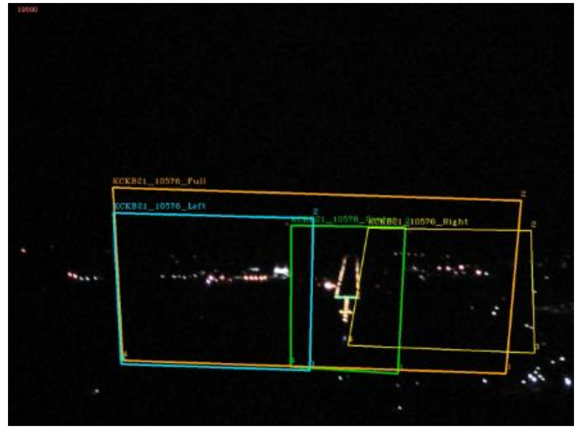

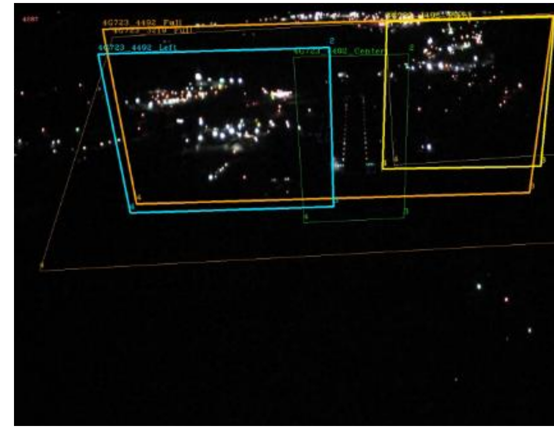

FRA 25L
Frankfurt DE

KCKB 21
Bridgeport WV

4G7 23
Fairmont WV

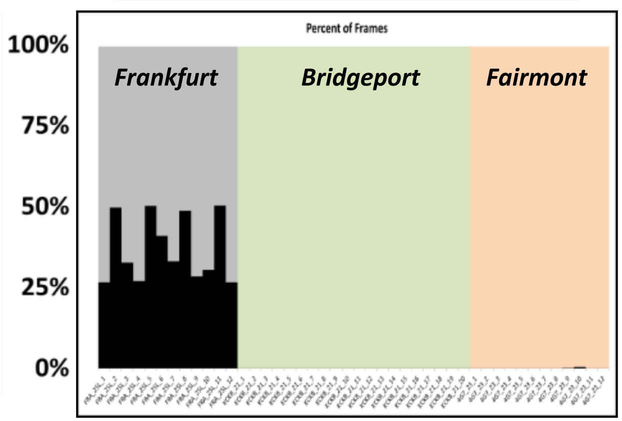

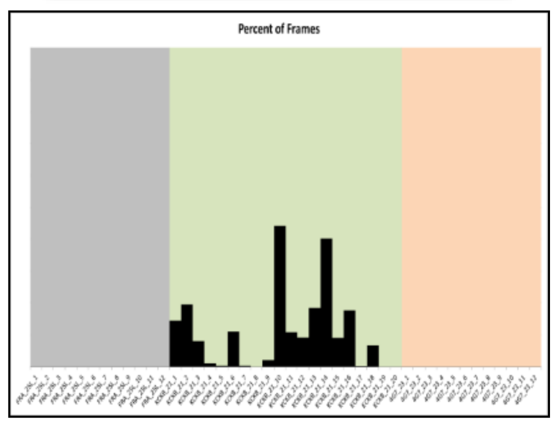

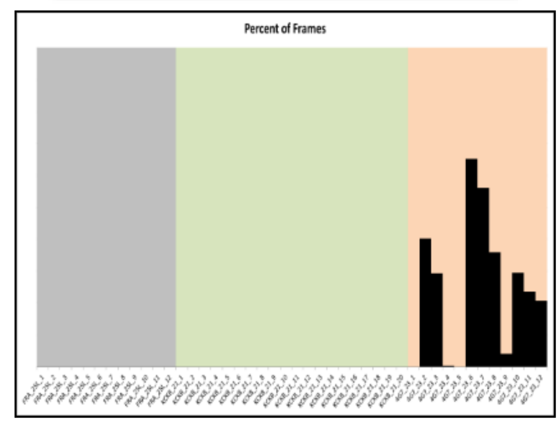

Hypothesis assumptions and predictions

For rigid objects and fixed scenes, current machine vision technology is capable of identifying imagery rapidly and with specificity over a modest range of camera viewpoints and scene illumination.

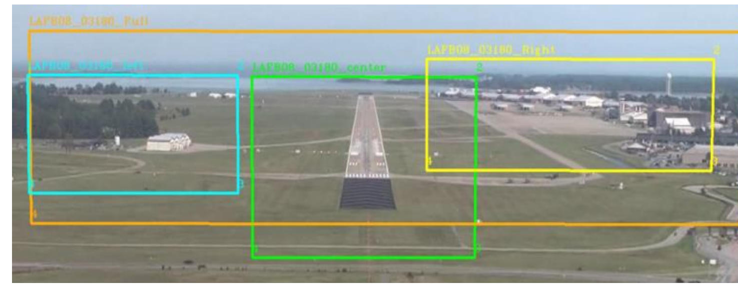

Source: NASA

Rigid objects	Runway video is recent – no major construction
Fixed scenes	No snow cover, full foliage
	Variation due to atmospheric turbidity - rain and fog in some approaches
Modest viewpoint change	Reject “cross-ways flyover” data
Modest illumination change	No dawn/dusk flights: two illumination variants (day and night).

Hypothesis assumptions and <u>predictions</u>

For rigid objects and fixed scenes, current machine vision technology is capable of identifying imagery rapidly and with specificity over a modest range of camera viewpoints and scene illumination.

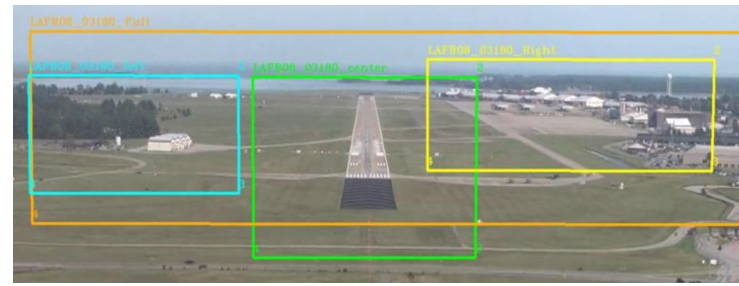

Source: NASA

Rapid	Yes: ~10fps for most tests
Specific	Yes: if the statistics are done correctly

"Stretch" prediction

Repeatable Only for

- same illumination (day vs. night)
- same sensor is used
- visibility was good

Data Prep

Raw videos

resolution:	visible - 1000x700 resolution or better infrared - 640x480
frame rate:	30 frames per second

Flights rejected

- perpendicular flyovers
- very distant views
- several from Google Earth
- several from simulators

19 videos, 9 locations

Approach Video	Runway	Band	Time	Weather	Way points	Refs/ waypoint
4G7_23_night	1		night	clear	3	4,4,4
EDDF_25L_night	2		night	clear	3	4,4,4
KAVC_01_fog_S	3 *		day	fog/rain	3	2,2,2
KAVC_01_fog_SE	3 *		day	fog/rain	2	1,1
KCID_27_LW	4 *	LW	day	fog/snow	6	4,4,4,4,4,4
KCID_27_SW	4 *	SW	day	fog/snow	6	4,4,4,4,4,4
KCID_27_VIS	4 *		day	fog/snow	8	4,4,4,4,4,4,4,4
KCID_27_LW_night	4 *	LW	night	clear	5	4,4,4,4,4
KCID_27_SW_night	4 *	SW	night	clear	4	4,4,4,4
KCID_27_VIS_night	4 *		night	clear	5	4,4,4,4,4
KCKB_21_night	5		night	clear	5	4,4,4,4,4
KLAS_25L_LW_night	6 *	LW	night	clear	4	4,4,4,4
KLAS_25L_SW_night	6 *	SW	night	clear	7	4,4,4,4,4,4,4
KLAS_25L_VIS_night	6 *		night	clear	6	4,4,4,4,4,4
KLFI_08_S	7 *		day	clear	3	4,4,4
KLFI_08_W	7 *		day	clear	4	4,4,4,4
KLFI_08_fog_NW	7 *		day	fog/rain	3	1,1,3
KLFI_26_rain	8		day	fog/rain	5	4,4,4,4,4
KPHF_25	9		day	clear	3	4,4,4
Total	9				85	321

Waypoints and Epochs

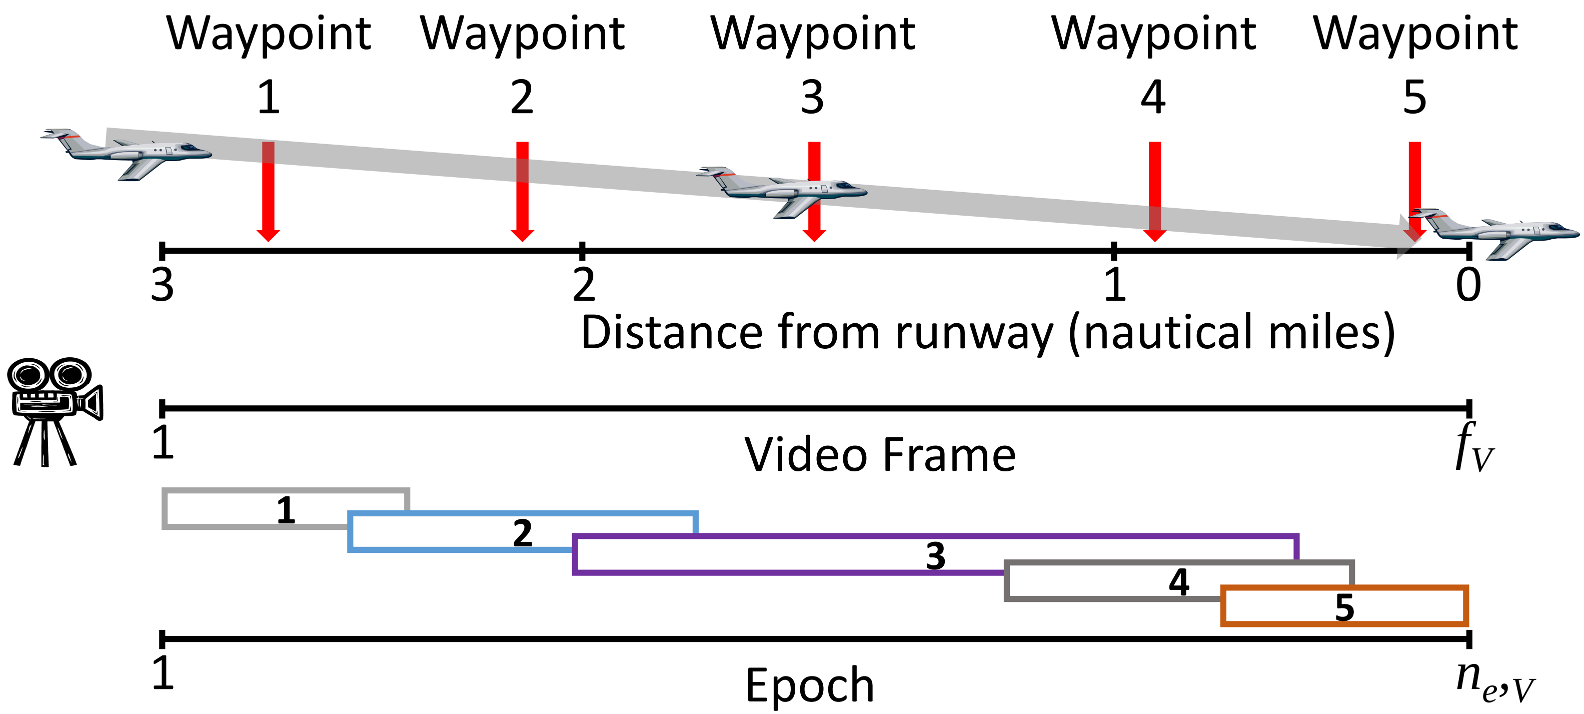

Source: NASA

Bayesian Probability Model

$$P(e|R_i) = \frac{P(R_i|e)P(e)}{P(R_i)} \qquad P(e) = \frac{1}{n_e}$$

Precompute:

$$P(R_i|e) = \frac{\sum_{f=1}^{f_e} M_{i,f}}{f_e}$$

$$P(R_i) = \frac{\sum_{e=1}^{n_e} P(R_i|e)}{n_e}$$

Compute on the fly:

$$P(e_t) = \left[1 - \prod_{i=1}^{k} \left(1 - P(e|R_i) \cdot M_{i,t}\right)\right]$$

KCID_27_SW_night

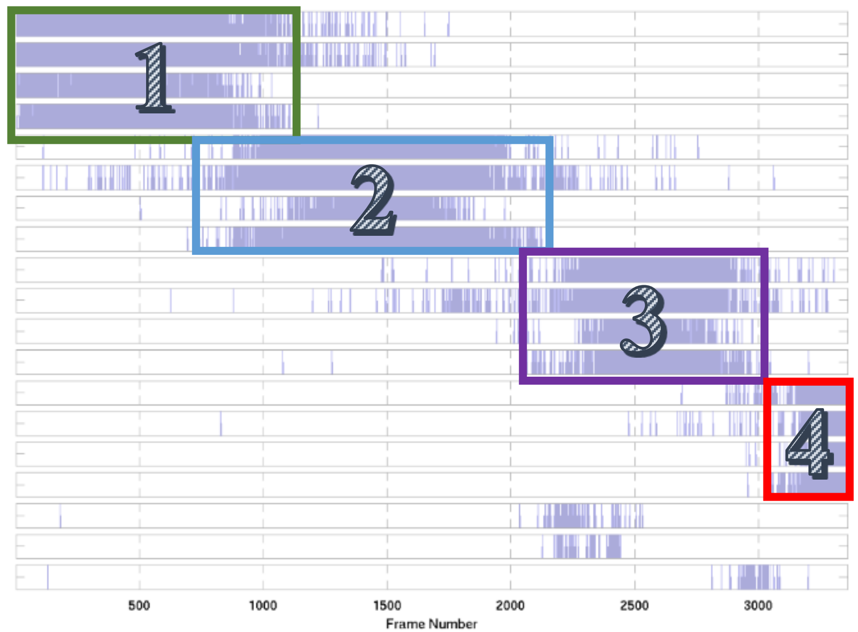

Epoch *e* is a proxy for location; location is unknown by default. "Unless one of the images I contain is matched in this frame, my probability is zero. Every match of one of the images that I contain increases my probability towards one."

Bayesian Probability Model

Source: NASA

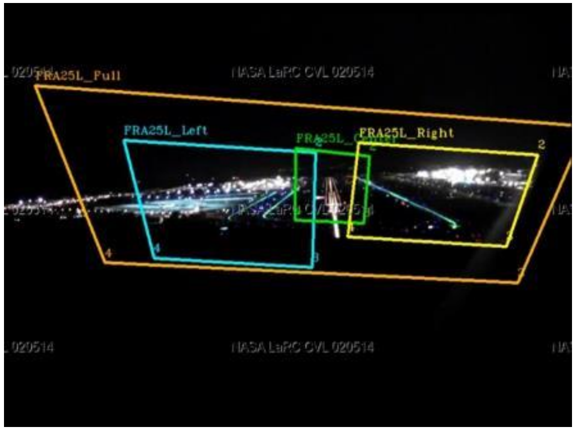

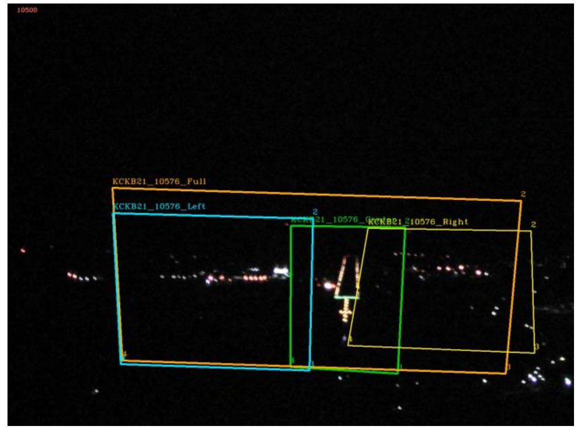

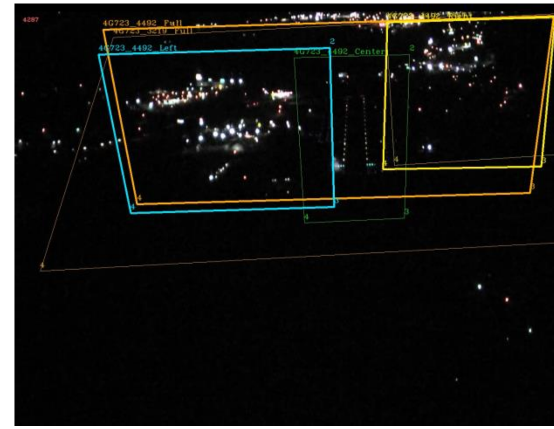

FRA 25L
Frankfurt DE

KCKB 21
Bridgeport WV

4G7 23
Fairmont WV

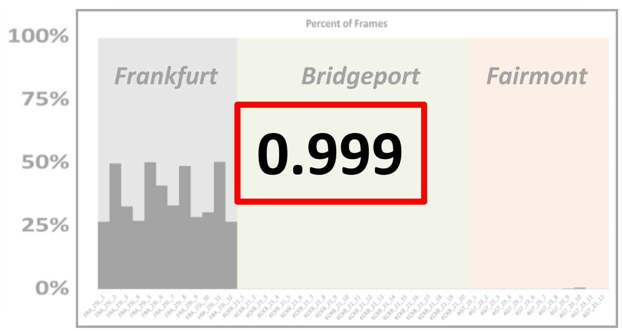

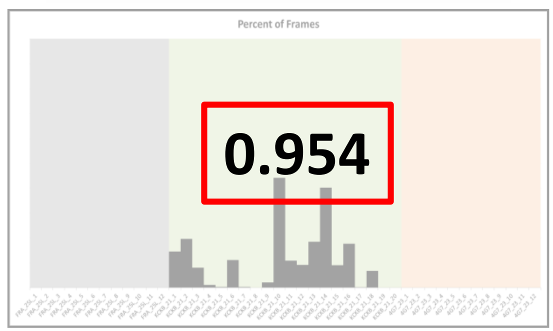

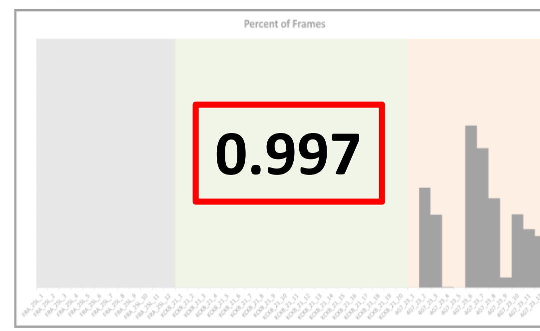

How to Represent Concisely?

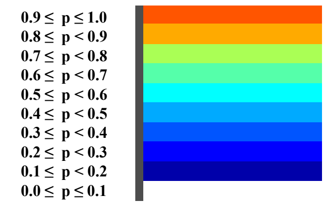

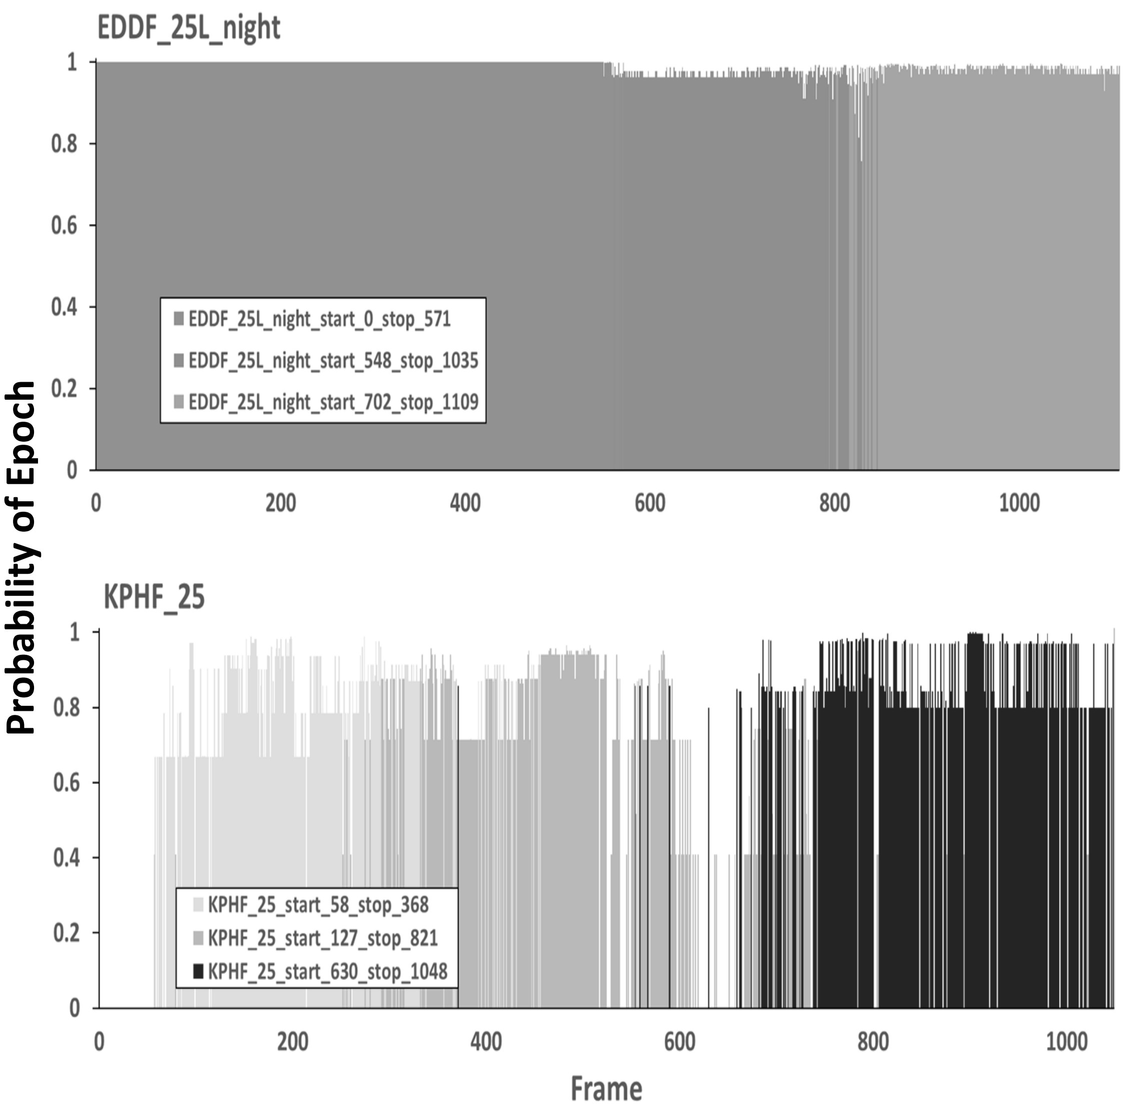

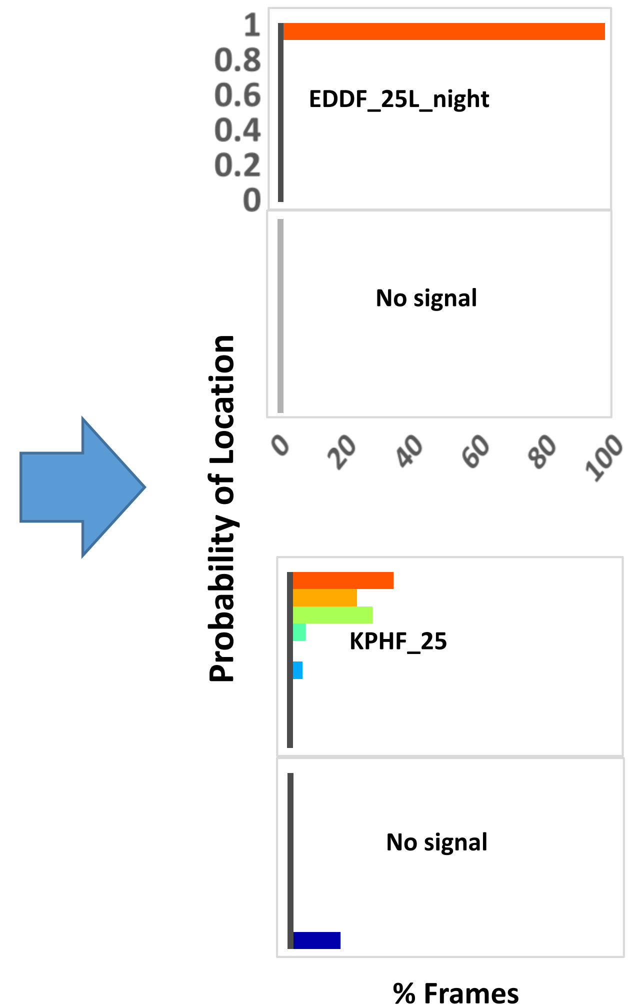

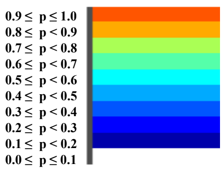

Results: Specificity

Estimated *vs.* True Location

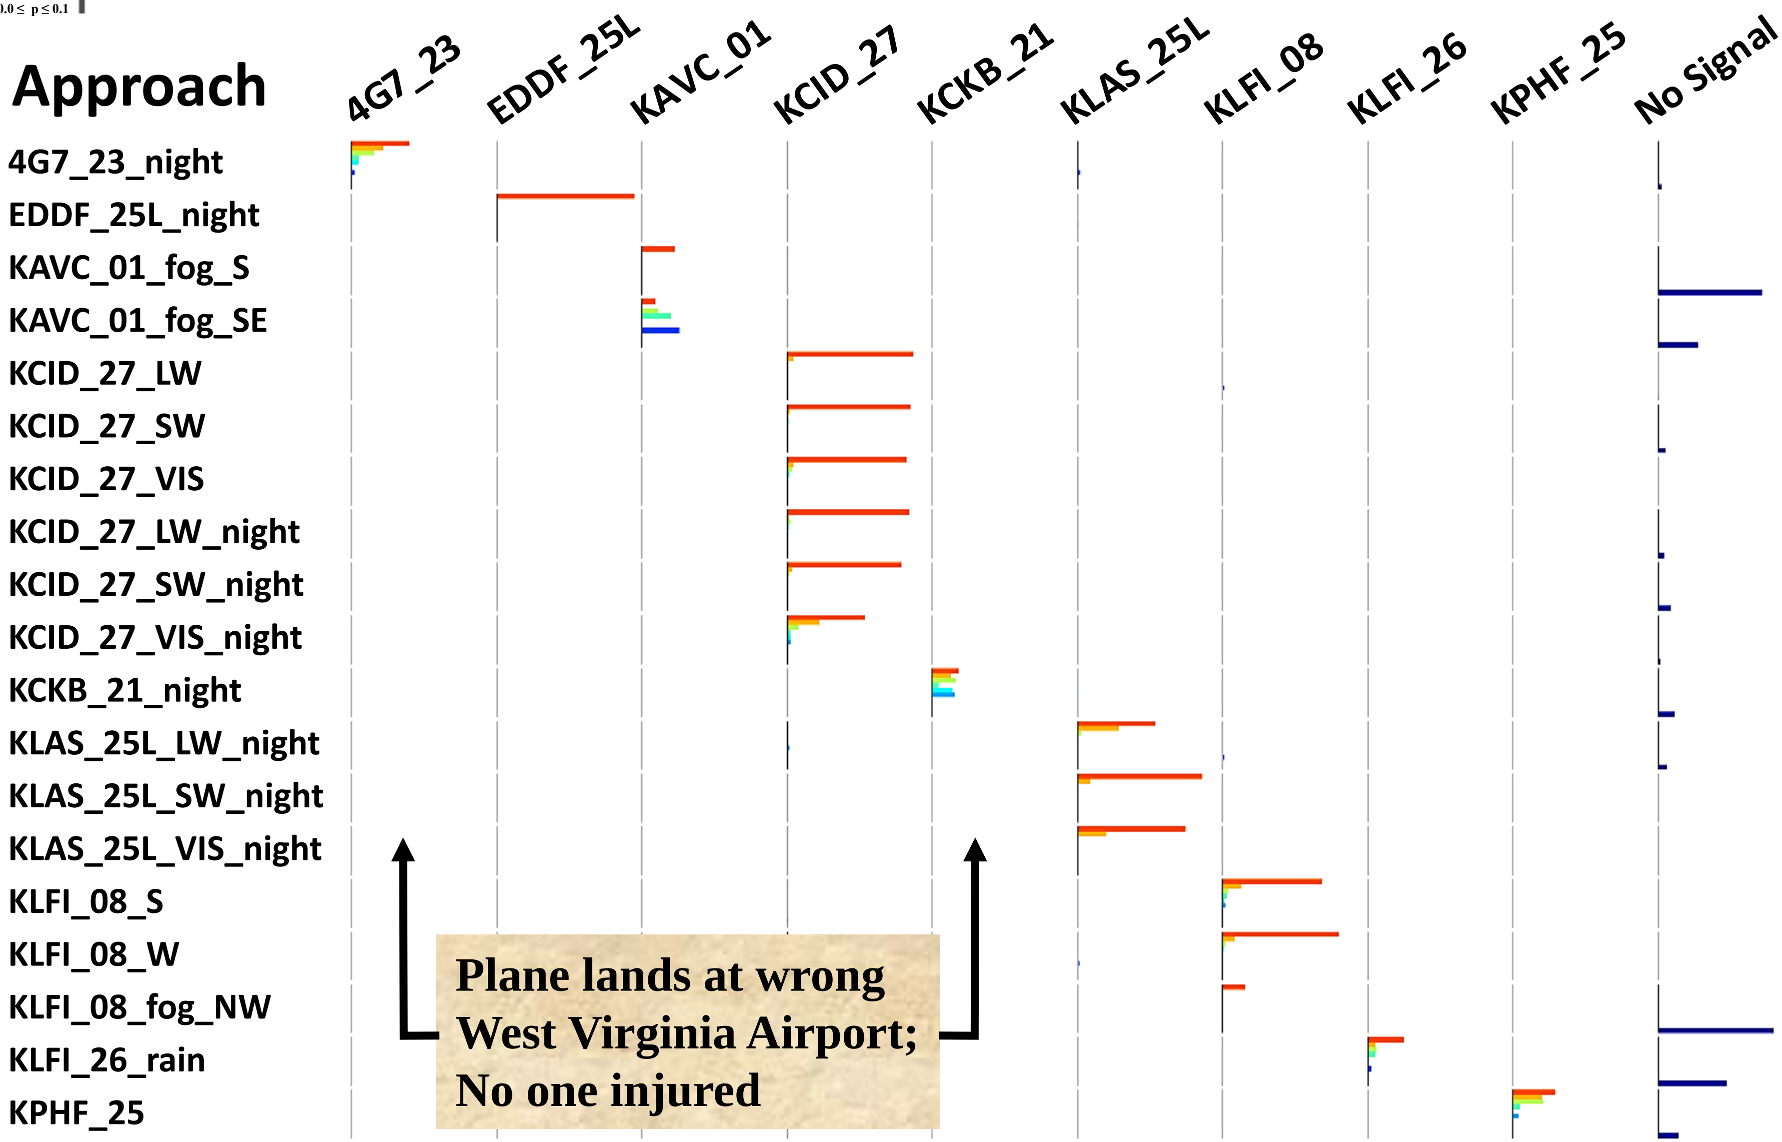

Results: Specificity (High Pass)

Estimated *vs.* True Location

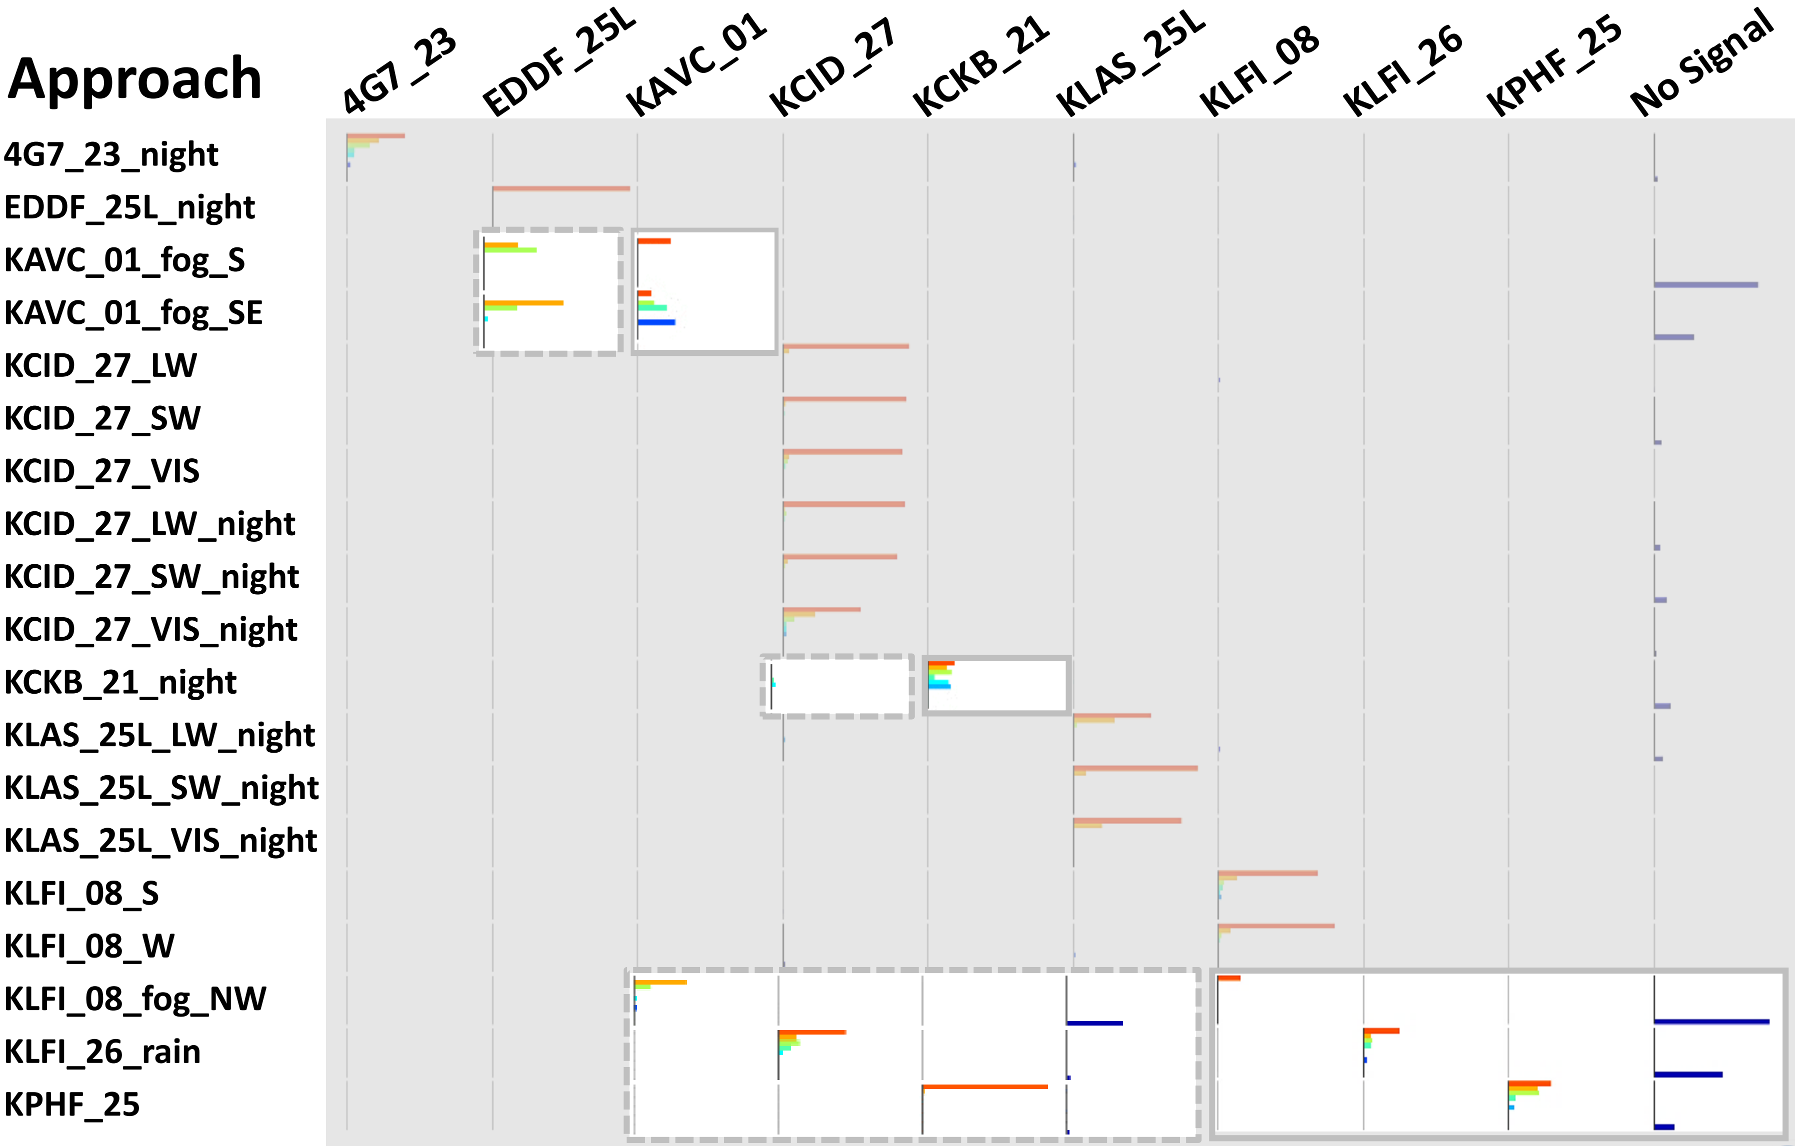

Results: Repeatability

Estimated *vs.* True Approach

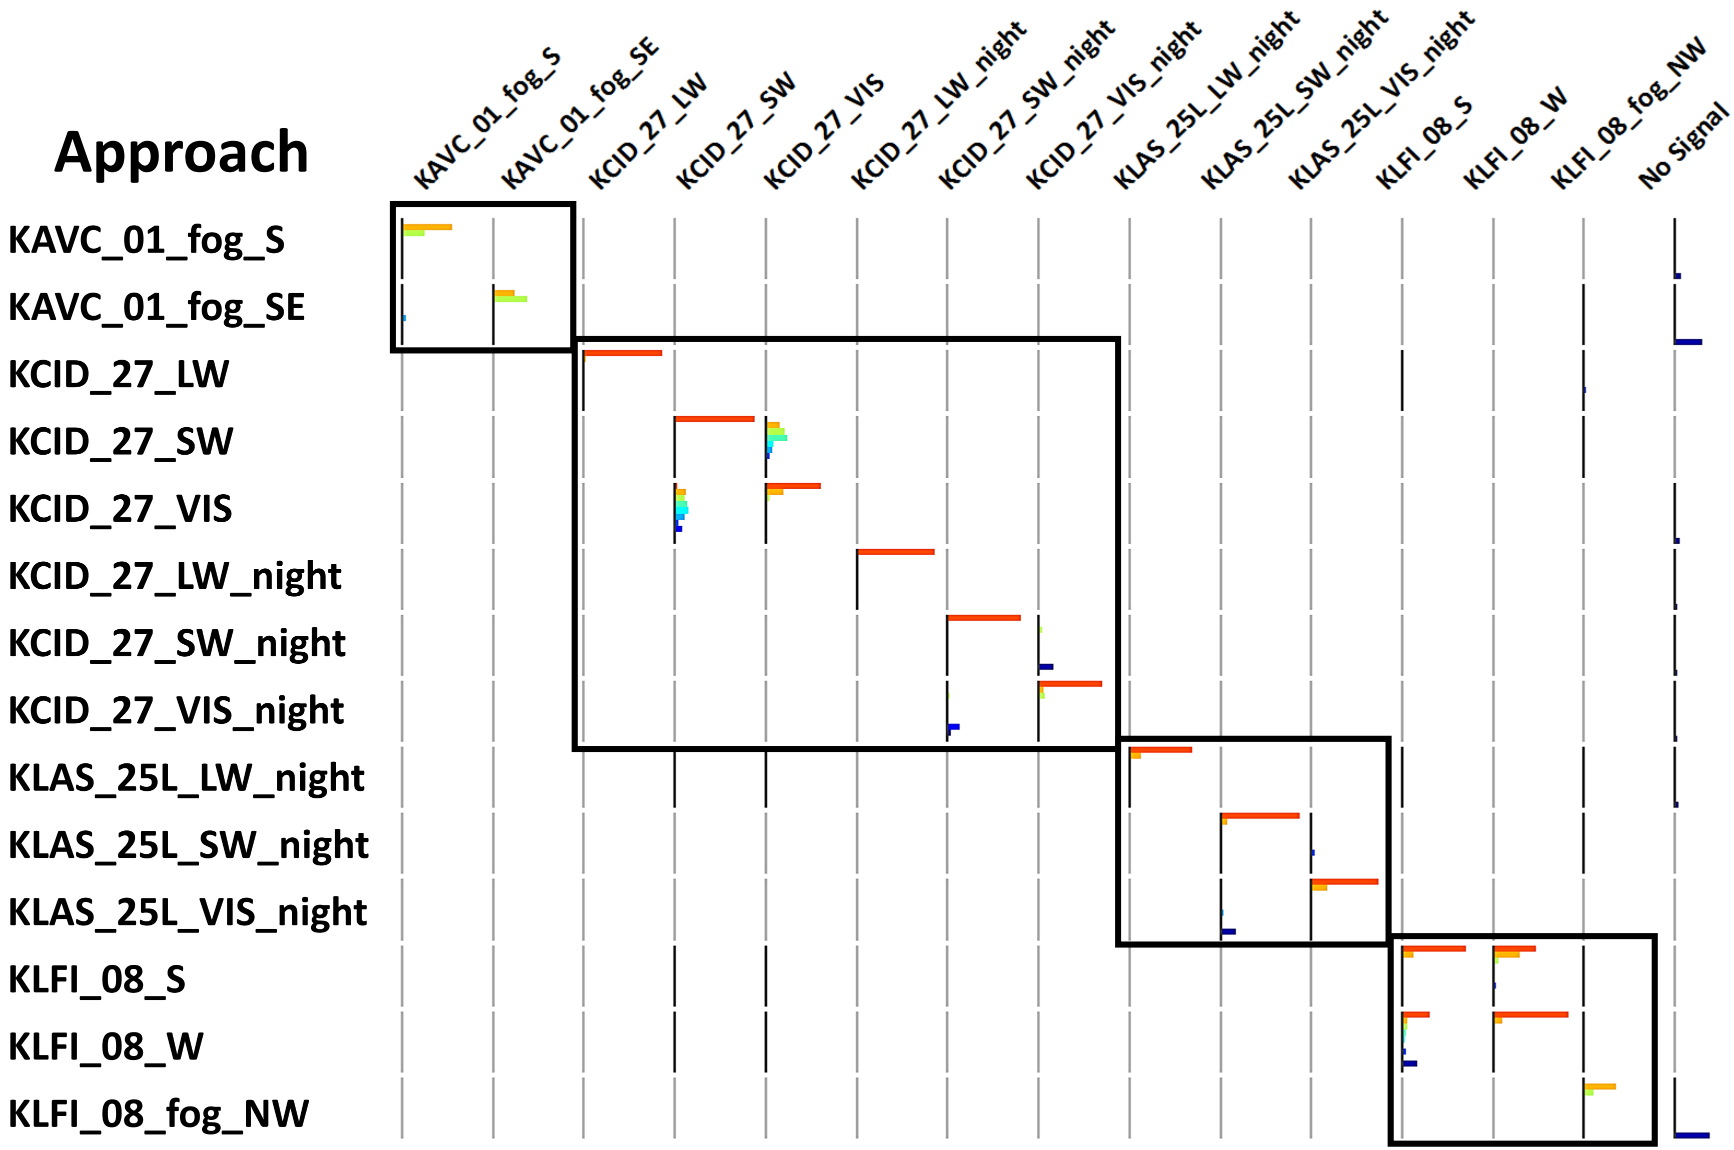

Cross-matching between visible and SWIR

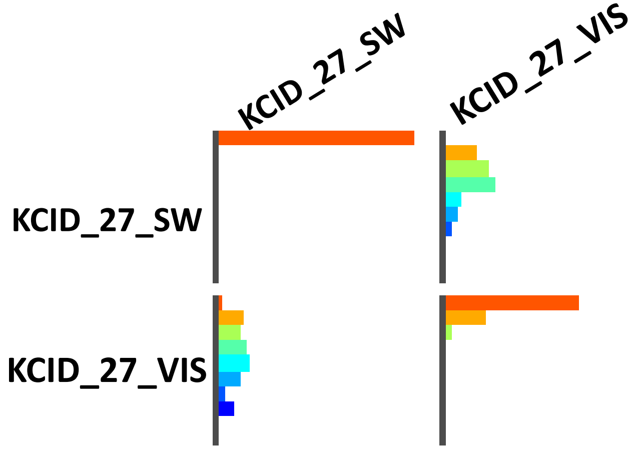

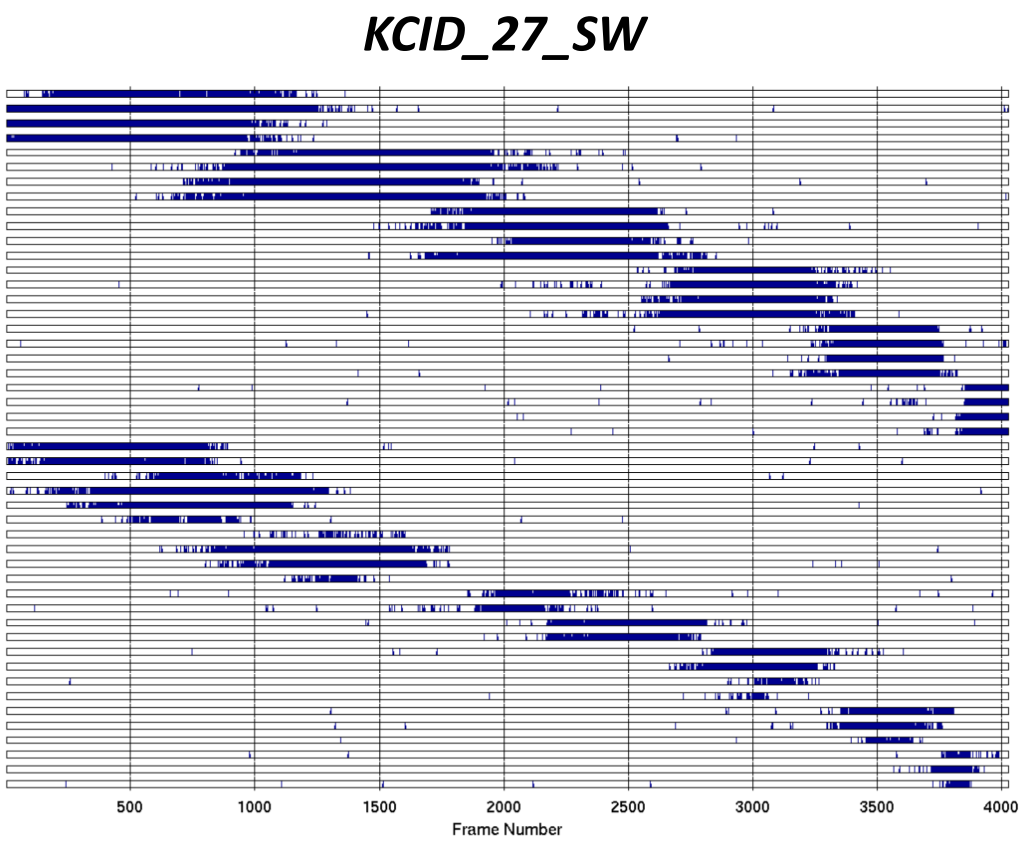

Conclusions

For rigid objects and fixed scenes, current machine vision technology is capable of identifying imagery rapidly and with specificity over a modest range of camera viewpoints and scene illumination.

Rapid ✓

Specific ✓ *

Repeatable ✓ *

* Robust if illumination (day vs. night) and sensor are the same.

* Both specificity and repeatability degraded in poor weather.

Surprise result: cross-sensor repeatability (visible <-> SWIR)

Geolocation via real-time comparison of cockpit video to a database is feasible, as long as

- the database contains imagery from the same time of day and
- the weather is clear at the time of the flight.

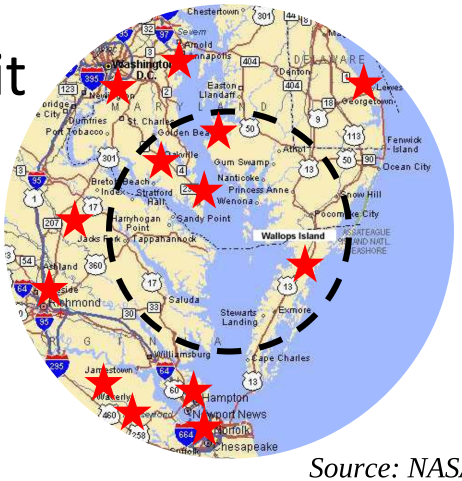

Source: NASA

Conclusions *Applications*

Geolocation validation and verification is already possible:

GPS – most of the approach

Inertial Guidance – all of the approach

Machine vision may (for clear flight conditions):

supplement these or

provide the same capability at lower cost,

though the certification cost is a significant barrier to adoption.

The limitations of machine vision geolocation in poor weather ensure that it *cannot be the sole backup* navigational technology.

Tests of its effectiveness with weather-penetrating sensors could overcome this limitation.

Anticipated over 15 years ago. Machine vision is now capable and compact enough to pursue this.

Backup

Literature Review (abbreviated)

Edge detection

Huertas et al. in 1990 used localized edge detection and thresholding to outline runways and create location-specific image templates for use in an expert system [12; see also 16].

Hough transform

Fleming and his collaborators reviewed the literature through 2004 [13] (including work that harnessed a key innovation, the Hough transform) and applied the Hough transform to runway imagery, using stereo ranging to estimate the airplane viewpoint and landing distance [14]. Independently and contemporaneously, Shang and Shi [15] took a similar approach, using monocular perspective analysis instead of stereo analysis to estimate the landing geometry.

Primitive learning

In his Master's thesis, Zongur [17] added a machine learning layer to previously applied techniques to recognize airports from orbital imagery.

Modern machine vision

Medioni and his colleagues at USC and Honeywell [18] used a new class of robust feature detector [2] and homographic perspective transformation to track runways in flight video; with image stabilization and image differencing they could determine if a runway was free of hazardous objects during the landing approach. Their application of the scale-invariant SIFT front end in 2009 represents a qualitative improvement in robustness in machine vision for aeronautics -- we use a performance-optimized variant (SURF [23]) in this study.

Next generation

A team at EPFL in 2014 applied methods at the current state of the art of machine vision to determine the boundaries of alternate landing sites such as agricultural fields in real time [20].

Epochs are containers of probability

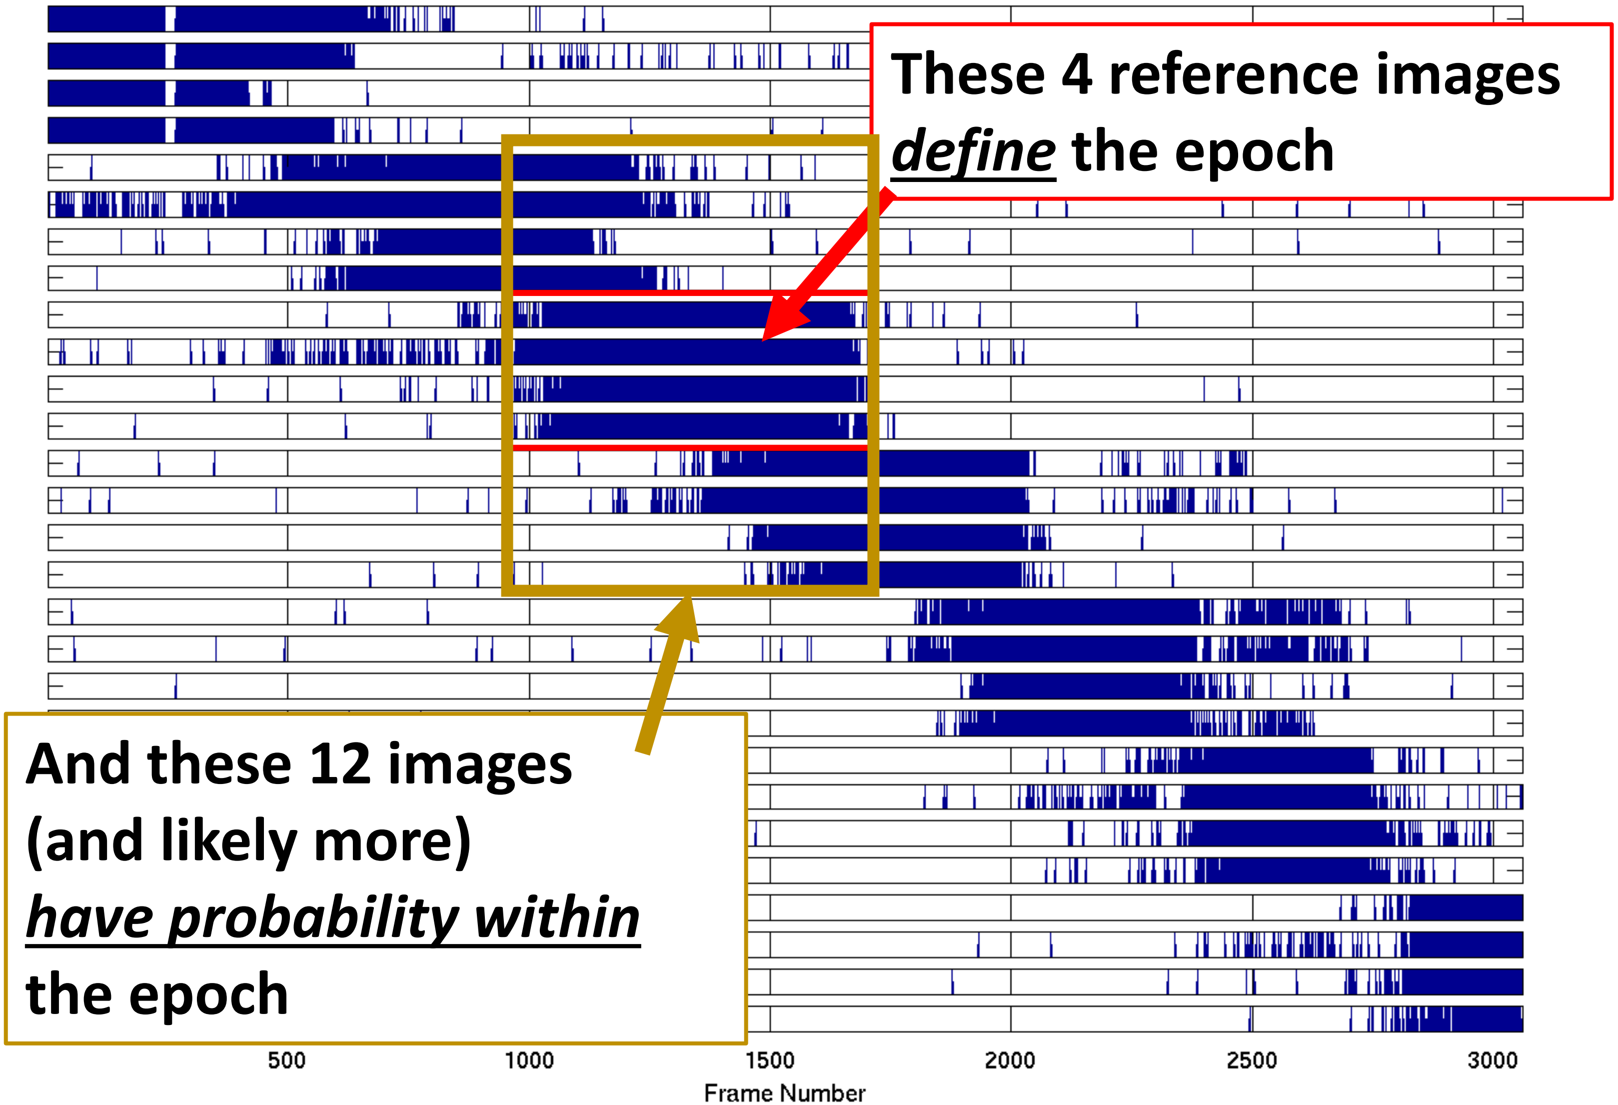

Conclusions

Aim of this study: explore the limitations of current machine vision technology, as applied to airborne geolocation, with realistic runway approach video taken in a variety of flight conditions. Our expectations:

- Rapid results
- Specific results
- Stability only for rigid objects
- Stability only for fixed scenes
- Stability only with modest change in illumination
- Stability only with modest change in viewpoint

Overall, the results met these expectations.

Conclusions *Speed, Stability*

Rapid results Linux computer with 16 CPUs and 4 GPUs -- 10 fps

Specific results As long as the flight conditions did not depart from the key constraints (rigid objects, fixed scenes, modest variation in illumination and viewpoint), location was determined with excellent specificity.

Rigid objects This constraint was maintained strictly throughout the study, in the sense that a clear runway or a patch of landscape viewed from the air is a fixed object.

For example, we did not test cases with obstacles on the runway, or cases in which heavy winds cause trees in the landscape to move noticeably.

Conclusions *Scene Variation*

The "fixed scene" constraint was varied in two ways.

1. In clear conditions we allowed minor changes within the scene: the moving vehicles. They had no discernable effect on the results.
2. A change in air turbidity is in essence a scene change. Four video approaches in fog/rain. Specificity was superb even with this violation of the fixed scene constraint. However, the constancy of specificity throughout the flight, was degraded. Contrast:
   - EDDF_25L_night -- strong location nearly 100% of frames
   - KAVC_01_fog_S -- strong location only 20% of frames

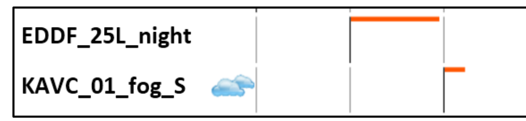

At first we suspected that constancy of location was degraded due to methodology, i.e., poor reference imagery choice or low waypoint count. But, *KAVC_01_fog_S* and *KAVC_01_fog_SE* seem to disprove this.

KAVC_01_fog_SE -- fewer waypoints and reference images but its location correct for over 50% of the approach (albeit with less than 90% confidence for much of that time).

Approach Video	Way points	Refs/ waypoint	Epochs	Refs defining epoch
KAVC_01_fog_S	3	2,2,2	3	1,1,2
KAVC_01_fog_SE	2	1,1	2	1,1

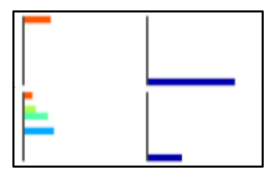

We did not use transmissometers to quantify visibility, and cannot identify a root cause.

Conclusions *Illumination change*

The "modest illumination change" constraint was maintained strictly throughout the study. Only mid-day and nighttime approach video was available.

Shadows: A change in shadows early or late in the day is an example of illumination change that violates the fixed scene assumption; more data is required to assess the fragility of machine vision due to shadows.

Diffuse vs. direct sunlight: What about, say, direct vs. diffuse sunlight on a foggy day?
We tested overcast conditions only. We assume that air turbidity is a stronger effect but cannot prove it with this data.

Conclusions *Modest viewpoint change*

The waypoint methodology ensures viewpoint constancy.

Simple conclusion: machine vision fails as expected when the viewpoint is changed greatly. For example, epoch locations from early in an approach have zero reported probability late in an approach.

Stronger conclusion: runway recognition with machine vision is effective with modest viewpoint changes.

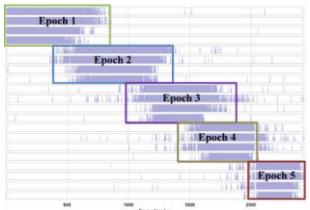

Example - KMEM_09_IR_night. What is remarkable is the time span of the positive result: reference images from a single 33 millisecond time sample (a single frame) produce a correct result for 15-30 seconds.

Location Confidence *vs.* Time

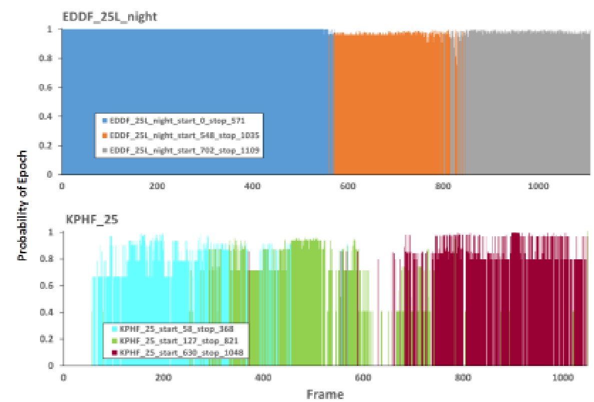

Important exception:

a) the continuity of reported location is less variable for a large aircraft (Boeing 747, at top) than a small aircraft (Cessna, at bottom)

b) the camera viewpoint in recordings from a small airplane is prone to sudden changes in pitch, yaw or roll – presumably due to wind shifts during the flight.

Since everything is recomputed for each frame, this does not arise from compute latency. Looking at the original video for *KPHF_25* and other Cessna recordings, we observed that the image sensor does not respond instantly to a sudden viewpoint change. Streaking, smearing and other frame readout artifacts are evident with each jerk. We conclude that a sensor with a higher frame rate, a faster pixel readout response, or frame-shielded design is needed to eliminate this effect. This sensor constraint should be observed in General Aviation and in small UAV applications.

Conclusions *Repeatability*

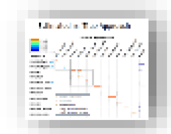

Clear daylight conditions: Machine vision can repeatedly locate a runway.

Clear night conditions: We expect that repeatability is possible but can’t prove it.

Fog or rain: Violate the fixed scene constraint and wipe out repeatability, even between two foggy approaches.

Across sensors and across time of day: Need more data to answer this.

Only simulator data was available to test across sensors (visible/SWIR) and across time of day

Because of these shortcomings, we cannot say that repeatability is beyond the capabilities of machine vision as the sensor or time of day is changed.

Simulator Pluses:

- Fidelity is sufficient to produce high specificity even with 640x480 resolution.
- Free from the sudden viewpoint changes as observed in the Cessna flights.
- Geodetics are superb – landscape features are identical across sensor types and times of day.

Simulator Minus:

Landscape features are not photorealistic.

- Nighttime light placement seems arbitrary as perspective increases toward the vanishing point.
- Ground features such as buildings are generally “flattened and painted” onto the ground.

Conclusions *Applications*

Geolocation validation and verification is already possible:

GPS – most of the approach
Inertial Guidance – all of the approach

Machine vision may

supplement these, or

provide the same capability (for some flight conditions) at lower cost,

though the effort to certify machine vision to the level of inertial technology is a significant barrier to adoption.

With visible and SWIR sensor inputs, the limitations of machine vision geolocation in poor weather ensure that it <u>cannot be the sole backup</u> navigational technology.

Tests of its effectiveness with weather-penetrating sensor inputs are required to overcome this limitation. Anticipated over 15 years ago.

Conclusions *Operational feasibility*

An aircraft with

- an onboard image sensor,
- modest computing power, and
- a database of reference images and probabilities

can geolocate as long as visibility is clear and the database

- includes imagery from various of times of day,
- covers the geographical areas likely to be encountered, and
- contains imagery of the same sensor type as the onboard sensor.

computation time will increase linearly with database size

We expect that the imagery database must also:

1. include seasonal variations such as snow cover, wet vs. dry pavement, and full vs. sparse foliage, and
2. be reasonably current. Minor scene changes did not degrade results, and we expect that similarly minor changes (traffic cones, commercial signage, and cell towers) will not do so either. Major changes, such as new roadways and buildings will violate the fixed scene assumption.