Daniel B. Ramey
James H. Smith
Lockheed Engineering and Management Services Company, Inc.
Houston, TX 77258

Simulation of Errors in a Landsat Based Crop Estimation System*

The simulator combines the use of empirical observations, theoretical probability distributions, and the reconstruction of archived weather patterns to provide a lifelike imitation of any of a large class of estimation procedures

Abstract: *Agricultural crop estimation has long been one of the primary applications of remotely sensed data. Development of advanced procedures for estimating crop production is a slow process, because the final evaluation of a procedure usually must wait until official agricultural reports are published sometime after the crops are harvested. By creating a generalized computer model of a Landsat-based estimation system, the authors have provided a tool that can answer many questions about proposed technologies prior to their implementation. The simulator described in this paper provides significant information beforehand by locating design flaws, aiding in the selection among competing component technologies, and providing iterative feedback for "fine-tuning" a procedure. The first part of the paper describes the simulation system (AgSim), and the second part gives an example of its use in a study of the effect of changing the Landsat orbit from an eighteen-day repeat coverage cycle to a sixteen-day cycle.*

Introduction

One of the earliest practical applications of remote sensing was in the estimation of domestic and foreign crop production. A major portion of the data collected by the Landsat series of polar orbiting satellites has been acquired over the world's grain producing regions, and Federal programs such as LACIE (1976-1978) and AGRISTARS (1980-1983) were established to promote research into developing efficient, economical crop estimation techniques (MacDonald, 1980; Erickson, 1982). In the past few years, much research has been dedicated to developing highly automated systems in which the data are collected, analyzed, and processed to final estimates quickly and economically, with minimal human interaction.

The development and testing of new estimation technologies normally requires that proposed methodologies be evaluated in a large number of situations, over diverse geographic regions and over a number of years. The expense of such a developmental process is accordingly very high. This paper describes a simulation tool designed to answer many important questions about new estimation technologies before large-scale implementation. This tool can be used to diagnose and correct design flaws in weeks rather than years.

The Agricultural Information System Simulator (AgSim) is an interactive computer program which models each of the major steps required to estimate a region's crop production. It allows the scientist to study the sensitivity of his final estimate to a variety of climatological and agroeconomical factors, and allows parameters of an estimation procedure to be fine-tuned in advance of its general application.

In the next section, we discuss estimation systems in general, with particular emphasis on potential sources of error in the ultimate estimate of production. As each source of error is explored, we briefly explain the method by which it is modelled and

---

* Presented paper, 17th International Symposium on Remote Sensing of Environment, Ann Arbor, Michigan, 9-13 May 1983.

Photogrammetric Engineering and Remote Sensing,
Vol. 50, No. 12, December 1984, pp. 1707-1712.

0099-1112/84/5012-1707$02.25/0

incorporated into the simulation. A description of the general design of AgSim is omitted, because its operation by design mimics that of an estimation procedure like the one described below. The paper concludes with an example of how the simulator was successfully used to assess the effects of changing the orbit of Landsat so that a point on the Earth's surface is visible to the satellite's scanner once every 16 days, instead of every 18. These intervals are in fact characteristic of, respectively, Landsat 4 and the earlier Landsats 1, 2, and 3.

SOURCES OF ERROR IN A LANDSAT-DERIVED ESTIMATE

This section presents a simplified system for collecting and analyzing remotely sensed crop data, and for making regional agricultural estimates based on these data. This system is more or less generic, but is typical of the techniques developed in support of the AGRISTARS project (see, for example, Dennis, 1982). Most systems proposed for large-scale application are similar in at least some respects to the one described here, so that the sources of estimation error described below are characteristic of many of the procedures developed to date.

Because high-resolution scanners collect vast quantities of data on a daily basis, a sampling scheme is usually employed to reduce the amount of data that must be analyzed. Sampling units are allocated to the region of interest, based on *a priori* knowledge of the area's agricultural, economic, and/or geographical characteristics. Data are collected for each sampling unit at each satellite pass through the growing season, then some procedure is employed which produces an estimate of the proportion(s) of the crop or crops of interest in each individual sampling unit. These sampling unit estimates are aggregated to produce large-area acreage estimates, which, when combined with yield predictions, provide the desired estimates of total production.

AgSim addresses the major error sources in satellite systems with the goal of understanding not only how individual factors affect the final estimate, but also what effects their interactions with one another have. A list of the potential sources of error in our hypothetical system are provided in Table 1. These sources are explored in detail in the following discussion.

SAMPLING AND RELATED ERRORS

Sampling and related errors are typical of agricultural surveys in general. A major factor contributing to the variance (or uncertainty) of the final estimate is the variability of the sampling units themselves. In some regions, the proportion of a major crop in an area the size of a sampling unit might be almost constant, while in other regions the proportions in neighboring sampling units may vary greatly. To increase the sampling efficiency of a fixed-size sample, a region is sub-divided into strata which are more homogenous with respect to crop proportions than is the entire region, and the sample units are apportioned among the strata according to some rule. A poor stratification for a region causes an unnecessarily variable final estimate, as does an improper allocation of the sampling units to the strata.

TABLE 1. SOURCES OF ERROR IN SATELLITE-BASED CROP ESTIMATE

Source	Contribution to Bias and Variance of Large-Scale Production Estimate
Sample Design	+ Stratification + Segment Allocation + Sampling Variance
Data Collection	+ Nonresponse due to: —Cloud Cover —Data Quality —Miscellaneous causes
Segment Proportion Estimation	+ Classification Errors + Discrimination Errors + Within-Segment Sampling
Aggregation to Regional Production Estimate	+ Ratioing to Replace Missing Data + Yield Estimation Errors

Because the sample design is an input parameter to AgSim, any sources of error arising from an incorrect or inefficient design are automatically incorporated into the simulation. Sampling unit to sampling unit variability is incorporated by modelling the distribution of sampling unit crop proportions over the population of all sample proportions. The required input parameters, mean stratum crop proportions, and sampling unit variance are then used to construct a probability density function from which a sampling unit proportion is drawn. In the case of a single crop of interest, the proportions are generated using the well known Beta probability density model. In the case of multiple crop estimation, a multivariate generalization of the Beta density known as the Dirichlet density is employed (Wilks, 1966). A more complete treatment of the generation of sampling unit proportions can be found in Ramey *et al.* (1982).

SEGMENT PROPORTION ESTIMATION ERRORS

Segment proportion estimation errors occur when estimates are calculated for each individual sampling unit. Some of the causes include inaccurate designation of field boundaries, incorrect labeling of the fields, and sampling errors, similar to those described above, which occur within the sampling units themselves. The simulator assumes that the following relationships hold among true sampling unit proportion (P), estimated sampling unit proportion (P'), and error (e):

$$P' = P + e$$

$$E(e) = aP + b.$$

In the above equation, $E(e)$ represents the expected value of e, b represents the (constant) bias of the estimate P', and the coefficient a allows the estimation error to exhibit a linear correlation with the true (or, as the case may be, simulated) sampling unit proportion P. Estimation techniques frequently tend to overestimate in sampling units of low true proportion, and to underestimate in sampling units of high true proportion (in other words, the coefficient a is usually negative).

Generation of the simulated sampling unit proportion estimates is very similar to the process of generating the "true" sampling unit proportions. Again, the Beta (Dirichlet) distribution is used to model the hypothetical population of sampling unit crop proportions, but this time the mean is the value P generated above, rather than the mean stratum crop proportion. The variance of the error e for each stratum is required input.

NONRESPONSE AND PARTIAL RESPONSE

Almost any type of sample survey is vulnerable to potential problems caused by lost observations. The term nonresponse, when applied to a sampling unit in our estimation scheme, means simply that what, if any, data collected for the unit during the growing season could not support the estimation of any of the crop proportions in the sampling unit. A unit can also suffer from partial response with certain estimation techniques. This occurs when only a partial estimate of the sampling unit's proportions is possible; e.g., instead of individual estimates for wheat and barley, the data representing a sampling unit might permit only a combined small grains estimate. All sampling units not suffering from nonresponse or partial response are said to exhibit full response. Note that we use the terms nonresponse, partial response, and full response with regard to sampling unit proportion estimates only, not with the loss of a single day's coverage of the sampling unit.

Most state-of-the-art crop proportion estimators require at least two reasonably unobstructed views of a scene, and further, these observations must lie within certain time "windows," determined by expected crop growth stage. The timing of data acquisitions determines whether the sampling unit will exhibit full, partial, or nonresponse. We refer to the list of calendar dates upon which data of acceptable quality are acquired for a sampling unit as that sampling unit's acquisition history. Once the temporal acquisition requirements for a particular estimation procedure are established, realistic patterns of full, partial, and non-response can be simulated by generating an acquisition history for each sampling unit in the study.

The major factor influencing the acquisition history of a sampling unit, aside from the orbital characterization of the satellite, is of course cloud cover. Realistic cloud cover patterns are simulated by using archived meteorological observations. While any source of cloud cover proportions could be used to drive the simulator, the two sources which are "on line" consist of a nine-year archive of ground observations for the U.S. Midwest, and a four year archive of data collected by AVHRR instruments on board NOAA polar-orbiting weather satellites. Both sources of data meet the requirement of providing observations at times close to that of the actual satellite pass (usually within one hour). A very brief description of the acquisition history simulator is given in the next paragraph. The interested reader is referred to Smith *et al.* (1982, 1983) for more detail. Also, see Bean (1981), Greaves (1971), and O. E. Smith (1979) for details of related cloud cover models.

The first step in simulating an acquisition history for a sampling unit is to determine the dates when the site would be in view of the orbiting scanner. Coverage by Landsat normally occurs every 18 days (16 for Landsat 4), but some areas are situated so that they receive coverage twice every 18 (16) days. Next, a year is chosen for its meteorological archive, and a grid is superimposed over the area under study. Then for any date on which a sampling unit would receive coverage, cloud cover from that data (in the year selected) is retrieved from the archive, and the observed data are used to fill in the grid (a bivariate interpolation routine is used to smooth the surface created and to fill in any "holes"). Note that cloud cover itself is not simulated, but it is rather reconstructed from archived data through the use of a number of empirically and theoretically derived relationships (Smith, 1982).

Now the surface created for a given day over an area represents the proportion of observed cloud cover at each point on the grid on the specified day. In particular, an estimate of the proportion of cloud cover at each sampling unit's location at the time of coverage is now available. The proportion of cloud cover is translated into the probability of obtaining a successful acquisition using an empirically derived function, and acquisition/no acquisition is simulated by comparing a generated random [0, 1] number with the acquisition probability. A random data dropout rate to account for miscellaneous factors other than cloud cover is user-selectable.

AGGREGATION ERRORS

Aggregation is the process of combining the individual sampling unit proportion estimates into a final estimate. One possible source of error at this stage is the loss of all sampling unit proportion estimates in a stratum. In that case, an estimate for the stratum can be approximated by ratioing, e.g., historically stratum *X* has had about half as much of crop *K* as has stratum *Y*, so we can use 0.5 times the current estimate for *Y* to arrive at an estimate for stratum *X*, when the latter cannot be estimated directly. Complex aggregation technologies (Feiveson, 1982), which use optimally derived weighting schemes, are designed to produce a final estimate with mimimal expected error. However, any contributions of the aggregation technique itself to the bias and variance of the final estimate would be difficult to predict in advance. The simulator is ideally suited for evaluating aggregation technologies before they are employed in practice. Similarly to the case of sampling schemes, the effects of the aggregation technique can be determined simply by applying the technique to the simulated sampling unit-level estimates.

A second cause of errors during aggregation is the variance and/or bias of the estimator of crop yield. The mean and variance of the yield estimate are input for each stratum (or more likely, for groups of strata, because yield estimates often are available only for large geographical areas). The model chosen for simulation of yield estimates is a simple normal probability distribution, and simulated yield estimates are drawn randomly from a population having parameters specified by the user.

EXAMPLE: A PRELIMINARY EVALUATION OF THE LANDSAT-4 ORBIT

Aside from carrying the advanced Thematic Mapper in addition to a Multi-Spectral Scanner, the 1983 launch of Landsat-4 marked a number of firsts for the remote sensing community. The Landsat-4 orbit differs markedly from that of its predecessors in that a point on the Earth's surface is scanned once every 16 days rather than once every 18 days. Because the width of the swath scanned has remained nearly the same (100 n.mi.), the amount of area which receives dual coverage (areas where the swaths of different days overlap) is consequently smaller with the new orbit. A question which naturally arises is whether this change in orbits affects the accuracy of large-area estimates, and if so, to what extent. The remainder of this paper describes a simulation study conducted by the authors prior to the launch of Landsat-4, designed to compare the two orbits from the perspective of the large area estimate.

The region chosen for the comparative study was the United States corn belt—Illinois, Indiana, and Iowa. An allocation of 180 sampling units from an earlier experiment provided the stratification and sampling unit locations upon which the simulation was based. Meteorological satellite data from 1975 were used to drive the cloud cover model, but yield data were not used in this study (sampling unit proportions were aggregated to estimates of total acreage only, not to estimates of production). Sampling unit proportions of corn and soybeans were generated based on USDA county level statistics for 1978, and from these proportions simulated estimates were created, using the observed variance of an experimental corn/soybean estimation procedure (Dennis, 1982). The requirements of this same procedure were applied to each sampling unit's simulated acquisition histories to determine which sampling units would be nonrespondents.

A total of 160 acquisition histories was simulated for each sampling unit. Eight orbital "slots" were chosen at two-day intervals for each of the two orbits (Landsat-4 versus predecessors), and ten replications of the experiment were performed for each slot/orbit combination. The term "slot" is used to distinguish between possible orbits of the scanner: Let us assume that the spacecraft only can follow a set path when an arbitrary point on Earth is being scanned, and that this point can only be scanned at a given time of day. Under these constraints, the satellite now is limited to exactly *N* possible orbits, or "slots," where *N* is the length in days of the repeat coverage cycle ($N = 16$ or 18 in our case). Finally, upon completion of the simulation, data on simulated acquisition rates, sampling unit response or nonresponse, and large-area estimates of soybean acreage were summarized and analyzed. The remainder of this section presents these results, their analysis, and the implications of this study.

Figure 1 (top) shows graphically the average number of simulated acquisitions per sampling unit as a function of the orbital slot. The overall means are 6.86 acquisitions per sampling unit for the 18-day orbit, and 6.85 for the 16-day orbit. The standard deviations are 0.52 and 0.29, respectively (the higher variability of the 18-day orbit is evident from the figure). As would be expected, the average acquisition rates for the two orbits are very close. The interesting phenomenon here is the disparity in the variances. At least for 1975, it would appear that the new orbit is less sensitive to the orbital slot effect than the previous orbit. The Analysis of Variance table (Table 2a) gives a statistical analysis of the data, with the orbit effect proving to be nonsignificant while the orbital slot effect was highly significant. The conclusions are thus that there is no indication that data acquisition rates will be affected by changing from an 18- to a 16-day orbit, but there is some evidence to suggest that the variability of acquisition rates may be smaller with the latter. In either case, the largest contributor to the variability of acquisition rates is the orbital slot that the satellite finds itself in at the start of the growing season.

The bottom half of Figure 1 depicts the simulated sampling unit processability rates, i.e., the proportion of sampling units which "respond" with estimates of their own crop proportions. Here, Landsat-4's orbit has a slight edge, with 32 percent process-

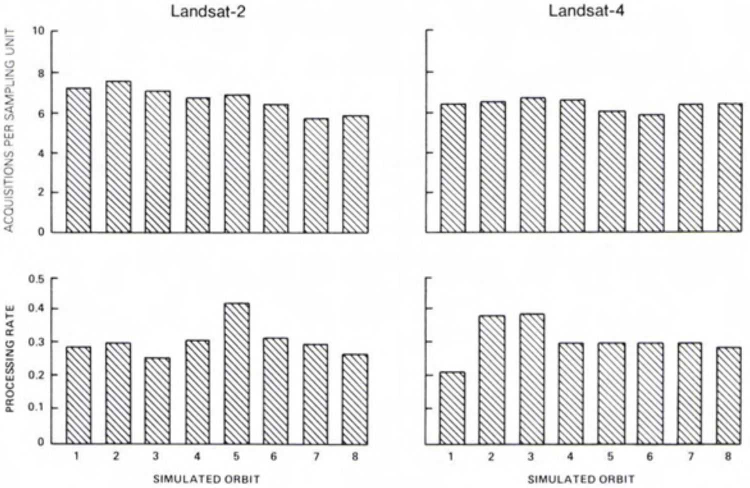

FIG. 1. Comparison of data acquisition and processability rates for Landsat 1-3 orbit with Landsat 4 orbit.

able compared with 30 percent for the predecessors. Surprisingly, the variances of the rates were very close, with Landsat-4 actually showing a slightly larger standard deviation (0.06 compared with 0.05). The analysis in Table 2b finds a significant orbit effect as well as a significant effect due to orbital slot. It appears that (at least in some cases) the final sample size is influenced strongly by the effectively random launch date of the satellite sometime in the past. While the limited scope of the present study hardly justified generalizing to all situations, one must be encouraged that the slight improvement in processing rate seen here might be characteristic of the 16-day orbit. Additionally, there is reason to expect that processability increases as the variance of the acquisition rate decreases, so the results from the acquisition-rate study above would seem to concur with the results here. In fact, the improvement in processing rate may be even greater than the 2 percent noted earlier. Of the 180 sampling units allocated, it was noted that a relatively higher proportion of the sample was located in the dual coverage regions of the Landsat-1/2/3 orbit than in those of the Landsat-4 orbit. When this factor is accounted for, the improvement rises to 3 percent absolute, or about a 10 percent expected increase in the final sample size.

TABLE 2. STATISTICAL TEST OF THE EFFECTS OF SATELLITE ORBIT AND ORBIT CYCLE/WEATHER CYCLE INTERACTION

(a) Average number of acquisition per segment

Source	DF	F value	P value	R^2
Model	87	15.81	.0001*	.95
Error	72			
Satellite orbit	1	.57	.45	
Orbit cycle/weather pattern interaction	14	94.87	.0001*	
Replications	72	.65	.97	

(b) Proportion of segments for which an estimate can be made

Source	DF	F value	P value	R^2
Model	87	6.31	.0001*	.88
Error	72			
Satellite orbit	1	5.77	.0189*	
Orbit cycle/weather cycle interaction	14	34.33	.0001*	
Replications	72	.87	.72	

(c) Bias of large-area estimates

Source	DF	F value	P value	R^2
Model	87	1.62	.47	.55
Error	72			
Satellite orbit	1	.05	.81	
Orbit cycle/weather cycle interaction	14	1.02	.42	
Replications	72	.92	.56	

Finally, an analysis was performed on the simulated estimates for total soybean acreage, and the results shown in Table 2c are all negative, that is, neither the orbit itself nor the orbital slot chosen had a significant effect on the error of the final estimate. Though the results are negative, they do point to an important conclusion: that the final regional estimate itself appears to be fairly insensitive to both the type of orbit and to orbital slot effects.

CONCLUSIONS

This paper has presented a simulation system which closely mimics the operation of a Landsat based crop forecasting technology. The simulator combines the use of empirical observations, theoretical probability distributions, and the reconstruction of archived weather patterns to provide a lifelike imitation of any of a large class of estimation procedures. One of the first practical applications for AgSim was described, and the analysis of the simulated results was presented. From this analysis, we see the sensitivity of both the old and new orbits to the orbital slot effect (though the new appears to be somewhat less sensitive than the old), the slight improvement in processability for the chosen estimation procedure afforded by the Landsat-4 orbit, and the relative insensitivity of the aggregated regional estimates to acquisition losses due to the orbit or to the weather. In summary, this experiment leads one to feel that the introduction of the 16-day orbit will have had no adverse effects on estimation technologies at the largest scale, and that perhaps more stability in the sample size will be realized.

Two final points should be emphasized: First, the high degree of realism achieved with AgSim did entail some cost—the historical archive of meteorological data that are presently machine readable is rather limited, so the results of simulations may not be generalizable to every situation. Second, the uses of the simulator are much broader than would be indicated by the example. Another AgSim study tested aggregated acreage estimates for a large region of the Soviet Union to see if local weather peculiarities would cause consistent bias in the large scale estimate over three growing seasons. Other possible applications might compare one, two, and three satellite systems, study the sensitivity of the final estimate to variations in the crop mix in the various strata, or evaluate the feasability of applying an existing corn/soybeans procedure in areas of South America.

REFERENCES

Bean, S. J., and P. N. Somerville, (1981). Some New Worldwide Cloud Cover Models, *J. Appl. Meteor. 20*, 223-228.

Dennis, T. B., R. B. Cate, C. V. Nazare, M. M. Smyrski, and T. C. Baker, 1982. SSG-4: An Automated Spring Small Grains Proportion Estimator, *Proceedings of the Eighth International Symposium on Machine Processing of Remotely-Sensed Data*, LARS-Purdue University, West Lafayette In.

Erickson, J., J. Dragg, R. Bizzell, and M. C. Trichel, 1982. Research Advances in Satellite-Aided Crop Forecasting, *Proceedings of the International Society for Photogrammetry and Remote Sensing*, Commission VII, International Symposium, Toulouse, France.

Feiveson, A. H., 1978. Weighted Aggregation, *Proceedings of the LACIE Symposium*, NASA Johnson Space Center, Houston TX, p. 1029-1037.

Greaves, J. R., D. B. Speigler, and J. H. Willard, 1971. *Developing a Global Cloud Cover Model for Simulating Earth-Viewing Space Missions* NASA Contractor Rep. CR-61345, 141 p.

MacDonald, R. B., and F. G. Hall, 1980. Global Crop Forecasting, *Science*, *208*, 670-9.

Ramey, D. B., and J. H. Smith, 1983. The Agricultural Information System Simulator: an Overview and an Application, *Proceedings of the Seventeenth International Symposium on Remote Sensing of the Environment*, Environmental Research Institute of Michigan, Ann Arbor, Mi.

Smith, J. H., 1983. Percentile Matching—An Alternative Method of Curve Fitting, *Proceedings of the Eighth Annual SAS User's Group International*, New Orleans, La, pp. 653-657.

Smith, J. H., J. T. Malin, C. C. Lin, and M. Dvorin, 1982. Acquisition History Simulation for Evaluation of Landsat-Based Crop Inventory Systems, *Proceedings of the Eighth International Symposium on Machine Processing of Remotely-Sensed Data*, LARS-Purdue University, West Lafayette In.

Smith, O. E., and P. N. Somerville, 1979. Worldwide Cloud Cover Model, *Fourth NASA Weather and Climate Program Scientific Review* NASA Goddard Space Flight Center, Greenbelt MD, pp. 315-318.

Wilks, S. S., 1966. *Mathematical Statistics*, John Wiley and Sons, Inc., New York, 644 p.

(Received 13 May 1983; revised and accepted 14 July 1984)

CIPA Slides on Architectural Photogrammetry

We have available for loan two series of slides prepared by the Laboratory of Photogrammetry, National Technical University, Athens, Greece. The first series deals with principles, equipment, and methods. The second series is titled "Application and Examples." These series are available on a free basis for two weeks to any member of the American Society of Photogrammetry.

Requests should be sent to:

US-ICAP
P.O. Box 705 (Uleta Branch)
Miami, FL 33164