Autonomous global sky surveillance with real-time robotic follow-up:
Night Sky Awareness through Thinking Telescopes Technology

**W. Thomas Vestrand, Heath Davis, James Wren, Przemek Wozniak, Ben Norman,
Robert R. White, Jeff Bloch, and Edward Fenimore**
Los Alamos National Laboratory

Barry Hogge, Moriba Jah, and Richard Rast
Air Force Research Laboratory

Abstract

We discuss the development of prototypes for a global grid of advanced "thinking" sky sentinels and robotic follow-up telescopes that observe the full night sky to provide real-time monitoring of the night sky by autonomously recognizing anomalous behavior, selecting targets for detailed investigation, and making real-time, follow-up observations. The layered, fault-tolerant, network uses relatively inexpensive robotic EO sensors to provide persistent autonomous monitoring and real-time anomaly detection to enable rapid recognition and a swift response to transients as they emerge. This T3 global EO grid avoids the limitations imposed by geography and weather to provide persistent monitoring of the night sky.

1. INTRODUCTION

The human brain, through a process we loosely call thinking, integrates data collection, pattern recognition, object classification, and memory together to obtain a higher understanding of what action needs to be taken and promptly takes action to respond to a threat or an opportunity. However, identification of ephemeral and/or subtle signatures of important anomalies in the huge data streams generated by night sky monitoring is too important to be left entirely to human analysts. Humans do not have the attention span, memory, or reaction time required to monitor huge volumes of data, recognize the important variations, and promptly respond with follow-up observations. To be effective, next generation systems will have to be autonomous. They must have a continuously evolving knowledge of normal behavior and be capable of recognizing subtle anomalies in a torrent of data. They must respond in real time by formulating queries and priorities, by commanding follow-up observations and by learning to optimize the response. They will therefore have to be thinking observational systems.

The technological building blocks for constructing robotic "thinking" observational systems exist: autonomous data collection, robotic hardware control, database construction and query, pattern recognition, object classification, and other forms of knowledge extraction. The goal of our Thinking Telescopes program is to merge those technologies to construct a full, end-to-end, globally distributed search engine for important events in the night sky. The integration of information science with robotic instrumentation will allow the monitoring and response instruments to feed observations back to a hierarchical decision engine, pose new questions for interrogation, conduct triage, and optimize the network hardware configuration for real-time knowledge extraction. Our key scientific goal is the construction of a global network capable of accurately localizing optical transients as short as a minute and making a series of interrogating follow-up observations while they are still bright. This global grid of sky sentinels and rapid response telescopes will observe the full night sky searching for optical transients, simultaneously monitor more than 10^7 celestial persistent sources and autonomously recognize anomalous behavior, select targets for detailed interrogation, and make real-time, follow-up observations—all without human intervention.

2. RAPTOR HARDWARE

At the core of our first generation prototype is a biologically inspired distributed sensor system that employs two arrays of telescopes, separated by 38 kilometers, to stereoscopically monitor the sky [1]. Each RAPTOR (*RAP*id *T*elescopes for *O*ptical *R*esponse) telescope array simultaneously images the same 1300 square-degree field with a wide-field imager and a central 4 square-degree with a narrow-field "fovea" imager (see Figure 1). The real-time software pipeline instantly analyzes images from RAPTOR A and B, potential candidates are identified, and the positions of any interesting transients without a measurable parallax are fed back to the mount controllers with instructions to point the fovea telescopes at the transient. The two fovea cameras then image the transient with higher spatial resolution and at a faster cadence to gather light curve information. Each fovea camera also images the transient through a different filter to provide color information. The RAPTOR A and B arrays therefore act as a binocular monitoring system employing closed loop feedback that autonomously identifies, generates alerts, and makes detailed follow-up observations of optical transients in real-time.

Previous attempts to search the sky for rapid optical transients from celestial objects with "monocular" systems failed because they were overwhelmed by false triggers. Those non-celestial false triggers are generated by a wide range of noise sources including—but not limited to—cosmic-ray hits, hot pixels, aircraft lights, image ghosts from bright stars, head-on meteors, and glints from space debris as well as satellites. To suppress false triggers, the "binocular" RAPTOR system uses two wide-field arrays (RAPTOR-A and RAPTOR-B) to stereoscopically view the same scene. The 38-kilometer separation of the telescope arrays yields a parallax shift of more than 220 arcseconds for non-celestial objects all the way out to the altitude of geostationary orbits at 36,000 kilometers. Our wide-field imagers have a single pixel resolution of 34 arc-seconds so any transient generated at distance at least out to six times geostationary has a detectable parallax. So, just like human vision, RAPTOR uses "binocular" imaging to distinguish distant objects from nearby objects and to suppress imaging faults found in a single "eye".
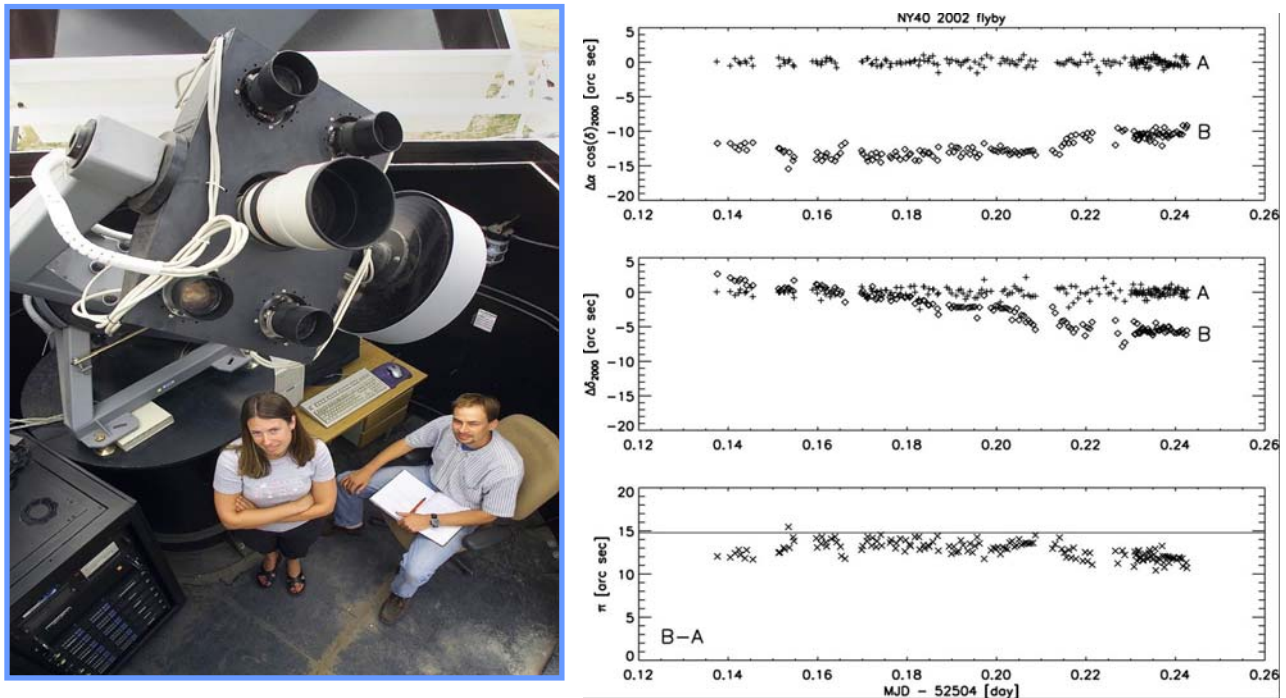
Figure 1 The first generation RAPTOR-A sky monitoring array is shown in the left-hand panel. As a demonstration of the stereoscopic ranging capability of the RAPTOR-A/B pair, we imaged asteroid 2002 NY-40 as it flew by at a distance of 1.3 times the distance to the moon. The measured parallax is plotted in the right-hand panel.

Our second generation wide-field monitoring array, RAPTOR-K, has unprecedented capabilities for persistent wide-field monitoring. It is composed of an array of 16 wide-field telescopes that simultaneously monitor nearly 1,000 sq-degrees (see Figure 2). The array is composed of sixteen 200mm Canon f1.8 lenses with 2Kx2K E2V back-illuminated CCDs at the focal plane. These systems reach a 3-sigma limiting magnitude of R=16th magnitude in 30 seconds. The entire array is mounted on a fast slewing mount that images the entire sky on a 5-minute circuit. This persistent monitoring array runs real-time software capable of identifying transients in less than 10 seconds and issuing alerts for real-time follow-up.

Figure 2. The second generation RAPTOR sky monitoring array. This array simultaneously monitors 10^3 sq-degrees with a sensitivity of R~16th magnitude in 30 seconds.

These sky monitoring arrays are supplemented by more powerful RAPTOR telescopes that make real-time follow-up observations. All of the telescopes are mounted on rapidly slewing mounts that are capable of pointing anywhere in the visible sky and beginning follow-up observations in less than 10 seconds. The RAPTOR-T(echnicolor) telescope array is composed of four high-end COTS optical tube assemblies that are co-aligned and co-located on a single mount (see Figure 3). Each of the four tube assemblies is a 0.41-meter f/10 telescope. Together these four co-aligned tube assemblies yield an effective aperture for RAPTOR-T of 0.82-meter. An advantage of this approach, beyond speed of development and lower cost, is the much more compact size of the telescope for a given focal ratio. This in turn translates into a lower moment of inertia and makes it much easier to achieve the extremely rapid slewing speeds. To make simultaneous multicolor observations of the transients identified by the sky monitors, each RAPTOR-T optical tube images the same field but through a different broad-band (clear-VRI) filter. The RAPTOR-T telescope array is now complete and operating at our Fenton Hill Observatory site. A limiting magnitude of R~18.5 in 10 seconds is achieved at our observatory site using a camera at the focal plane that employs a 1Kx1K Marconi back-illuminated CCD chip. All four cameras are synchronized in order to perform simultaneous multi-color imaging.

Figure 3. The RAPTOR-T multi-color follow-up telescope array (foreground) and the RAPTOR-K sky monitoring array. These fully autonomous robotic telescopes monitor the night sky for interesting targets and make real-time follow-up observations without "humans in the loop".

## 3. MACHINE LEARNING AND ANOMALY DETECTION

Even without "thinking", this telescope array would produce important science. But the true power of the system will be its ability to learn to recognize important variations in persistent sources, as well as identify transients, and command follow-up observations while the ephemeral anomalies are still present. Therefore a crucial and challenging aspect of our thinking robotic system is the development of tools that enable automated real-time processing and knowledge extraction from imaging data generated at a rate approaching a terabyte per night. Techniques have been developed for knowledge extraction both from images and time series data. But can one automate some of those techniques and integrate them with a robotic observing system to enable real-time follow-up observations? Speed, sensitivity, and the suppression of false positives are important.

Figure 4. Anomaly detection software together with the record of normal source variations will be used to identify any important variations and alert our more powerful telescopes for real-time follow-up observations. The displayed field is a piece of one of our real images (each white dot is a star) and shows only 2% of the area simultaneously monitored.

While our accelerated photometry/astrometry pipelines conduct much of the real-time processing for the autonomous telescope system, the understanding of most astronomical objects is too incomplete to predict the properties of important changes. Our approach is to create a record of the observed variations for every persistent source and to develop new algorithms for anomaly detection along with machine learning techniques to train the system to recognize important changes as they are happening. Various "real world" artifacts, from airplane lights and clouds to sensitivity variations for individual pixels and other non-celestial phenomena that cannot be predicted, will inevitably contaminate the observations. Automated identification of these unwanted artifacts is an essential aspect of the monitoring software. Rather than "hard wire" the anomaly pipeline in an often *ad hoc* manner, we are employing Machine Learning (ML) techniques to train a classifier to identify and filter these artifacts based on a database of examples. This enables the system to become "smarter" and to continuously improve the overall quality of the data collection and system response. This approach also allows the users to interact with the ML software and to construct queries such as "find more like this" or if you ever detect something like "this" notify me and make it a high priority for prompt follow-up observations. The system therefore acts as an evolving search engine continuously monitoring a dynamic database—the night sky.

Anomaly detection is an established branch of machine learning that is a fundamentally unsupervised learning method. Using unmarked data samples (from sub-samples of the archive itself), the anomaly detection algorithm provides a "simplest" specification of the data. Similar objects cluster in complex regions of n-dimensional parameter space. Objects in the various feature regions can therefore be flagged as artifacts unworthy of follow-up or as real objects with a priority ranked by the astronomer for follow-up. Data not described by this specification are treated as anomalies and identified for follow-up studies to determine its properties. Many anomalies are generated by instrumentation errors---automating their identification helps the system to "bootstrap" its knowledge and improve the response. But some of the anomalies are rare events that are difficult to find any other way.
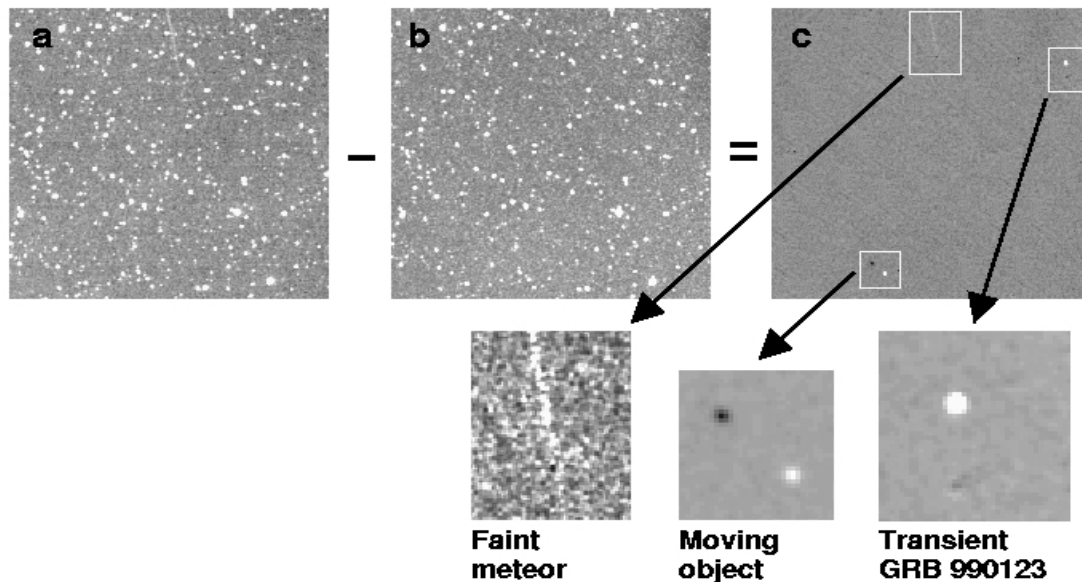

Figure 5. An example of anomaly detection and classification. Employing our Difference Imaging Analysis (DIA) software on real images collected with one of our sky monitoring telescopes at LANL shows examples of real transients.

4.   MEMORY AND CONTEXT

Persistent monitoring and construction of a baseline record of temporal variability is essential for distinguishing normal variations from anomalous variations. Without context information, the system cannot "think and learn", it can only apply static criteria to newly collected data in the analysis pipeline. Our system provides context through a Virtual Observatory (VO) that stores the measurements and learned experience (derived information and relations found by ML algorithms). This VO also acts as a platform for mining and knowledge extraction. To maximize adaptability of the anomaly detection software and scientific utility of the VO, we have constructed a powerful user interface that enables extensive data mining and integration with other databases.

The prototype of this VO that is publicly available is called SkyDOT (*Sky D*atabase for *O*bjects in the *T*ime Domain). SkyDOT is currently a *static* VO that contains a full year of our observations for ~10 million sources covering the full sky visible from Los Alamos (see figure 6). This VO is a useful tool for getting a quick start in the development of ML and anomaly detection algorithms as well as a reasonable temporal baseline to explore with new algorithms. This *static* version is now being expanded to be a *dynamic* Virtual Observatory for real-time astronomy with continuous updates and rapid verification of alert candidates generated by our thinking telescopes. This is a challenging goal and has required the development of new metadata standards for temporal databases and the development of astronomical extensions for database management and query languages.

Numerous high level tasks like classification, clustering and pattern recognition have to be passed to software agents. We are developing a framework for mining of the spatio-temporal data commonly found in sky and earth monitoring applications. This framework provides extensive cross-reference and searching capability based on the indexing algorithms: Hierarchical Triangular Mesh (HTM) and Healpix. We also use the Extensible Markup Language (XML) to define new data formats that are well suited for interoperability between VOs.
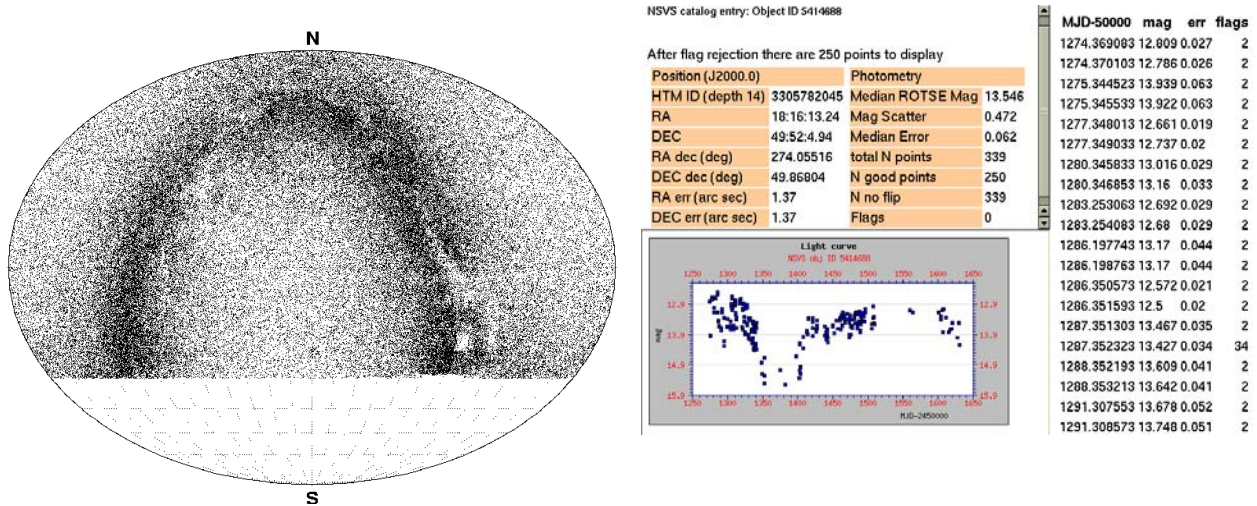
Figure 6. The oval on the left shows objects brighter than 11[th] magnitude (~500,000) detected from LANL plotted in an equal area projection. Each dot represents an object in SkyDOT with a temporal history typically composed of a few hundred plots covering a one-year baseline. The panels on the left show a sample record for one object. (See http://skydot.lanl.gov).

## 5. THE THINKING TELESCOPES NETWORK

The core goal of our Thinking Telescopes program is to fuse Information Science, Robotic Instrumentation, and Distributed Sensor Network technology to construct a full, end-to-end, globally distributed search engine for cosmic explosions. Figure 7 shows our proto-type Thinking Telescopes network as it currently exists at Los Alamos National Laboratory. A revolutionary aspect of this system is that it employs a coordinated network of intelligent software agents to conduct sky monitoring and transient recognition on a global scale. The integration of information science with robotic instrumentation allows the monitoring and response instruments to feed observations back to a hierarchical decision engine, pose new questions for interrogation, conduct triage, and optimize the network hardware configuration for real-time knowledge extraction. Our ultimate goal is construction of the world's first global network capable of accurately localizing optical transients as short as a minute and making a series of interrogating follow-up observations while they are still bright.

Figure 7. Components in the Thinking Telescopes network at Los Alamos National Laboratory.

One of the important capabilities found in the Thinking Telescopes network is the ability to trigger successive assets, increasing the type and number of systems used in any observation sequence based on the initial information gained from monitoring (see Figure 8). When the monitoring system detects a new object, interesting event, or an important change in the behavior of the persistent object, it will then bring to bear the next available resource in the chain to image the event. This secondary interrogation will then be further analyzed to verify the first observation. The schedule and scientific criteria will be reevaluated and then larger resources will be triggered for additional follow-up observations. This chain can be repeated any number of time as long as there are assets available.

Another important aspect of our Thinking Telescopes System is its ability to also reconfigure some assets as needed. To fit some science goals an adjustment to the actual optics may be necessary (e.g. changing filters or re-aligning OTAs) or adjusting observation criteria to support previous observations (e.g. switching to high cadence imaging). This flexibility in adjusting assets based on scientific goals is a key for any advanced astronomical system and an advanced distributed sensor network.

Figure 8. The existing Thinking Telescopes network and the timescales possible for finding and making follow-up observations of fast transients.

6. PATHFINDING SCIENCE

Since the discovery of prompt optical emission from gamma-ray bursts [3,4], our goal has been to make measurements while the gamma-ray burst (GRB) is still emitting gamma rays. The distribution of GRB durations is bimodal and suggests a class of short GRBs with durations less than ~2 seconds and there is growing evidence that short GRBs have different progenitor than long duration GRBs. Do short GRBs also have prompt emission that also extends down to optical wavelengths? Optical observations might be able to answer that question, but it is impossible to locate them and slew a narrow field telescope to observe them before the GRB over. One needs to continuously monitor a significant fraction of the area being monitored by the gamma ray detectors. The utility of persistent monitoring by RAPTOR was dramatically demonstrated recently when it detected bright optical emission from the "naked eye burst", GRB080319B, before it was localized by gamma-ray satellites [5].

There is evidence for the existence of celestial optical transients that cannot be found through sky monitoring by high-energy satellites. An exciting result from our RAPTOR observations of GRB 060206 was the fully autonomous

detection of high redshift explosive optical flaring without gamma-ray emission [6]. GRB 060206, which occurred at z=4.045, had a modest gamma-ray fluence (15-350 keV fluence~8×10^{-7} ergs/cm^2) and a gamma-ray temporal profile composed of a single, Gaussian-like, peak with duration T_{90}=7±2 seconds. However, nearly an hour after the gamma-ray pulse the emission from the optical counterpart underwent a spectacular re-brightening by ~1 mag. with a flux doubling timescale of about 4 minutes (see figure 9). The total R-band fluence received from GRB 060206 during this flare is 2.3×10^{-9} erg cm^{-2}. In the rest frame of the burst this corresponds to an isotropic equivalent energy release of E_{iso}~0.7×10^{50} erg in just a narrow UV band of $\lambda\approx130\pm22$ nm. This detection of bright, ~16th magnitude, "optically rich" flaring from a gamma-ray faint GRB at high redshift is a strong motivation for conducting untriggered searches for fast optical transients even with relatively small survey instruments. But, to search for optical transients, one needs not only a wide-field optical monitoring system but also one that can autonomously locate, in real time, celestial transients and command follow-up observations for transients with timescales as short as a minute.
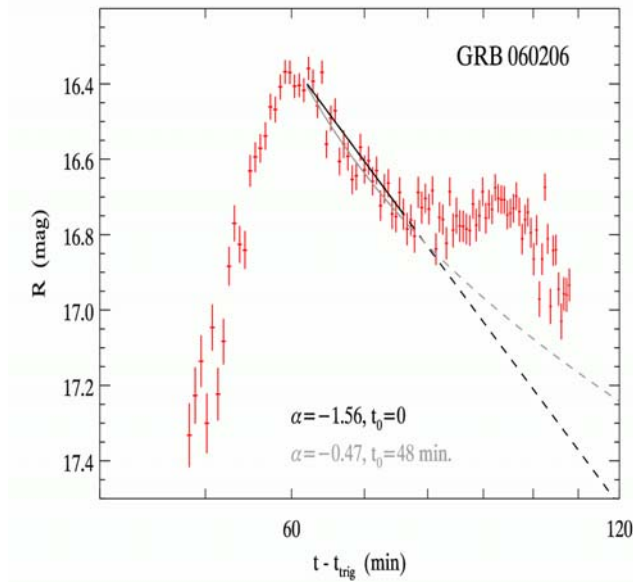
Figure 9. The light curve for the late time flaring measured by RAPTOR-S from GRB 060206. At the onset of the explosive re-brightening the flux doubled on a timescale of 4 minutes. The fading behavior of the flare is consistent with a power-law flux decay with index ~-1.

7. REFERENCES

1. Vestrand, W. T. et al., *SPIE* **4845**, 126, 2002.
2. Vestrand, W.T. et al., *Astron. Nach.*, **325**, 549, 2004.
3. Akerlof, C. et al., *Nature*, **398**, 400, 1999.
4. Vestrand, W.T. et al., *Nature*, **435**, 178, 2005.
5. Wozniak, P.R. et al. *Astrophys. J.*, in press, 2008.
6. Wozniak, P. R. et al., *Astroph. J.*, 642, L99, 2006.