# Commentary on: Monitoring and Diagnosis On-Board Software Module for Mars Driller (MODI) Authors: Jameux, Vitulli, Ribeiro, Fonseca, Santos and Barata

**Stephen F. Smith**
The Robotics Institute
Carnegie Mellon University
Pittsburgh PA 15213
sfs@cs.cmu.edu

This paper describes the MODI project, an ongoing effort to demonstrate the viability of an on-board software module for Monitoring and Diagnosis of remote systems, in particular those planned for deployment within the context of Martian exploration. The specific remote system of interest in the current prototype development effort is a drilling and sampling system that is expected to be onboard an early rover mission of the European Space Agency's (ESA's) Mars Exploration program.

As pointed out in the paper, the ability to autonomously monitor and diagnose various mechanical systems (together with the companion ability to replan and reschedule activities in response to detected problems) is truly an enabler of robotic planetary exploration. Given both the substantial delay in communications with earth-based ground stations and the inherent uncertainty of remote rover activities, the prospect of ground-based mission control (via telepresence) is problematic - at best, it will result in an extremely inefficient execution process, with corresponding degradation in scientific return. Efficient mission operations in this setting instead requires a division of responsibility, where ground-based decision-making takes a more strategic focus, and the mechanics of foreseeing task-level execution problems and managing task execution are handled locally (via onboard software systems). In light of this requirement, the project described in this paper is well motivated and important.

The paper takes what it calls a fuzzy, knowledge-based approach to developing a monitoring and diagnosis module for the Martian drill application. Under this approach, variables that describe various sensory inputs about aspects of the drill are modeled as fuzzy sets, and the important relationships between these variables/sets are expressed as rules. The paper summarizes the current prototype model, consisting of 3 "set points" (variables, such as the drill's RPM speed, whose values are imposed by humans and assumed to be non-varying), 7 drill-related parameters (variables for which nominal values and ranges are derived from data and monitored for alerting conditions), and 1 environment parameter characterizing terrain (to be incorporated in a future version of the prototype). A set of rules defined over these set points and parameters characterize the set of conditions under which an alert should be issued, with the fuzzy set membership functions associated with the variables in the rule determining the severity of the alarm. The domains of various drill parameters, the set of 36 distinct drilling scenarios for which monitoring rules were specified, and the set of monitoring rules were all developed through joint analysis between the authors and the company responsible for manufacturing the drill. The fuzzy set membership function of each drill variable for each scenario was derived from mean and standard deviation data using a machine learning process. Initial experiments designed to confirm the correctness of the inference process are reported and show the system's ability to appropriately detect alarm conditions of various severity.

Overall, the use of fuzzy logic would seem a natural approach to specifying a monitoring and diagnosis module for the drilling application, given the predisposition for natural variation (uncertainty) in sensor readings. The model developed in the paper seems to provide a promising initial solution, and no doubt this model will be enriched and refined as the project progresses. The projected capability to detect terrain type, for example, should greatly increase the robustness of the current prototype.

To its credit from the perspective of producing a demonstrable prototype, the paper takes a rather straightforward approach to model specification - i.e., identify all possible scenarios, construct specific membership functions for each variable for each scenario, and develop one or more rules for each scenario to delineate alarm conditions. However, this leads me to ask two broad questions.

- How amenable is the approach to accommodating more sophisticated types of monitoring situations? For example, one assumption that seems to be made in the current model is that the value of any given parameter is independently symptomatic of a problem but there is no combined effect. I could imagine situations where slightly abnormal values of a single variable might indicate a negligible problem whereas as several slightly abnormal values occurring simultaneously might indicate increased likelihood of a problem. Is this sort of rule expressible in the fuzzy knowledge-based framework?

  Another more advanced form of monitoring that might be enhanced by a fuzzy representation is that of detecting unfavorable trends (involving examination of variable values over multiple time steps). In the paper, the task of trend analysis is assumed to be responsibility of the ground station, and this certainly makes sense from an initial application perspective. However, it is interesting to think about the advantage that might be taken of a fuzzy representation of alarm severity to enable early detection of problems in an extended autonomy circumstance.

- How scalable and maintainable are fuzzy monitoring and diagnosis models? One of the practical myths of rule-based systems is the idea that one just focuses on encoding the knowledge and control takes care of itself. In practice, as rule sets get large, this becomes harder and harder to do and management of the rule set becomes a knowledge intensive activity in its own right. The paper's approach to automating the determination of fuzzy variable membership functions relieves some of the knowledge acquisition burden. It may also be the case that monitoring and diagnosis of mechanical devices of a robotic payload are sufficiently contained domains that scalability problems do not arise. However, it would nonetheless be interesting to contrast the fuzzy knowledge-based approach advocted in this paper with more model-based diagnosis approaches from a scalability/maintainability standpoint. One assumption made by this latter class of approaches is that physical systems are decomposable and system dynamics can be modularly specified through the construction and composition of subsystem models. I wonder if there is an analog to this kind of modularity (and hence scalability) in the fuzzy, knowledge-based approach, to support application to more complex types of remote systems.