Learning Locomotion

Proposer Information Pamphlet (PIP)

for

Broad Agency Announcement BAA 05-25

Defense Advanced Research Projects Agency
Information Processing Technology Office
3701 North Fairfax Drive
Arlington, VA 22203-1714

TABLE OF CONTENTS

1 PROGRAM OBJECTIVES

The Defense Advanced Research Projects Agency (DARPA) Information Processing Technology Office (IPTO) is soliciting proposals for a new program that applies learning technology to land locomotion in extreme terrain. The goal of the Learning Locomotion program is to develop a new generation of learning algorithms that enable traversal of large, irregular obstacles by unmanned vehicles.

Large, irregular obstacles such as urban rubble, rock fields, and fallen logs present minor challenges to dismounted forces, slowing but not stopping them. Thus, the obstacles create Slow-Go areas for dismounted forces. For today's small unmanned vehicles, these same obstacles cause No-Go areas, limiting effectiveness on the battlefield. Enabling future unmanned vehicles to traverse large, irregular obstacles will allow robots to better contribute to military operations.

Locomotion over extreme terrain requires deliberately planned, precisely coordinated movements. Like a hiker traversing a boulder field, the unmanned vehicle in extreme terrain will succeed not by flailing, but by meticulously sequencing its motions. The complexity of the planning and the degree of the required sensorimotor coordination presents significant challenges to the design and implementation of control systems. Handcrafting the control laws and parameters may not even be possible with reasonable effort.

Automatic learning offers a promising alternative. In the Learning Locomotion program, algorithms will be created that learn how to locomote based on the experience of a legged platform confronting extreme terrain. It is expected that the performance of these algorithms will far exceed the performance of handcrafted systems, creating a breakthrough in locomotion over extreme terrain. Further, it is expected that these algorithms will be broadly applicable to the class of "agile" ground vehicles.

2 PROGRAM DESCRIPTION

Figure 1 depicts the expected schedule for the program, which will proceed in two phases over approximately three years. Figure 1 shows the Learning Locomotion program Milestones, which call for traversals at speeds and over obstacles that relate to L, which represents platform leg length. The Milestones are:

- Milestone I: Learned system travels at 0.1 L per second, traverses obstacle of height 0.4 L
- Milestone II: Learned system travels 0.6 L per second, traverses obstacle of height 0.9 L

Meeting the milestones are criteria for all performing teams to continue to Phase II; as such, performers are competing against program metrics and not each other for down selects from Phase I to Phase II.

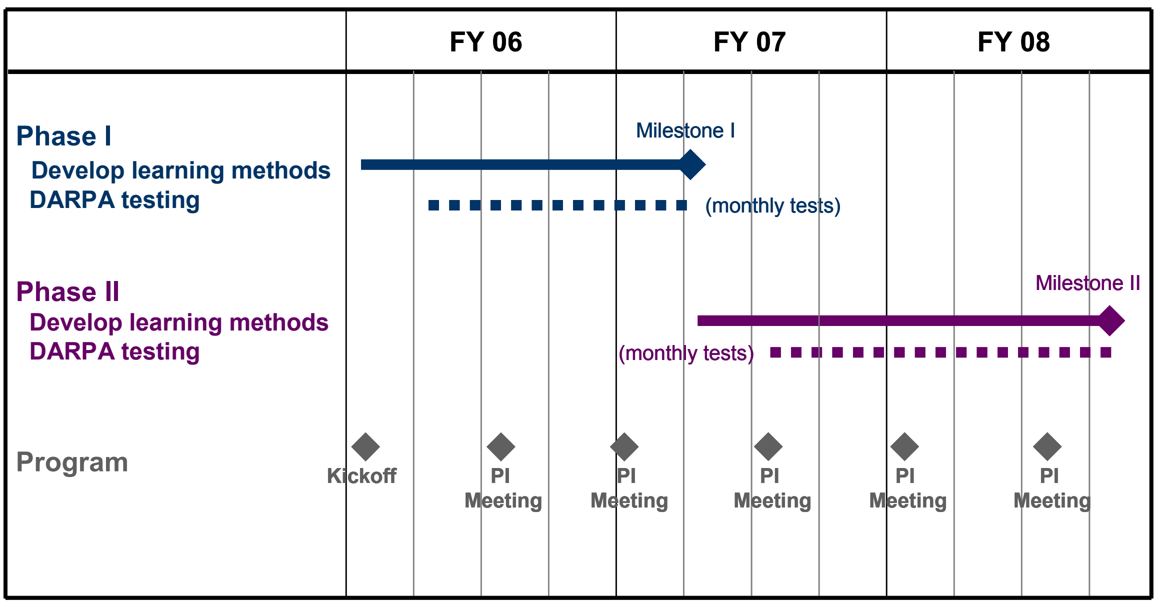

Let L represent leg length

Milestone I. Learned system travels at 0.1 L per second, traverses obstacle of height 0.4 L

Milestone II. Learned system travels 0.6 L per second, traverses obstacle of height 0.9 L

Figure 1. Current Program Schedule

Participating teams will receive one "Little Dog" locomotion platform and one (or possibly more) Terrain Board(s) as Government Furnished Equipment (GFE). The Little Dog robot is a custom-made four-legged platform approximately 10.5 inches in length and with a leg length of approximately 6 inches. Terrain Boards are transportable topographic boards, approximately 4 feet by 8 feet, which present a series of obstacles for the robot to traverse.

Teams will develop learning algorithms and test them by having Little Dog traverse courses on a Terrain Board. Performance will be measured by travel speed, and by the size of the largest obstacle traversable. Periodically, teams will upload their algorithms to a central facility that will conduct independent tests, using an identical Little Dog on a test Terrain Board.

Simple maintenance to the Little Dog platform and Terrain Boards will be performed by the performing teams. As a result, teams should have some basic minor mechanical/electrical in-house capability. Major repairs to Little Dog platforms and Terrain Boards will be performed by the Government team but at a cost to the performing team's budget and with the lost opportunities of practice and testing.

Table 1 details the performance expected by phase. To be allowed to continue into Phase II, performer teams will be required to provide software enabling Little Dog to achieve the Phase I performance metrics.

Table 1. Performance Metrics. The metrics are scaled for the Little Dog platform (L=6.3 in) and for a human with L=29.5 in.

Attribute	Units	Phase I	Phase II
Speed	L/sec	0.10	0.60
Little Dog	In/s	0.6	3.8
Human	In/s	2.9	17.7
Maximum Obstacle Height	L	0.40	0.90
Little Dog	In	2.5	5.7
Human	In	11.8	26.6

3 LOCOMOTION PLATFORM

3.1 LITTLE DOG

Each participating team will receive one Little Dog locomotion platform (Figure 2). A prototype of the Little Dog has been developed and is now undergoing testing.

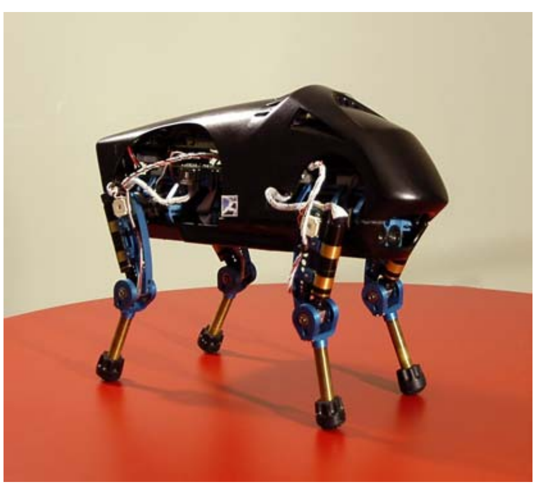

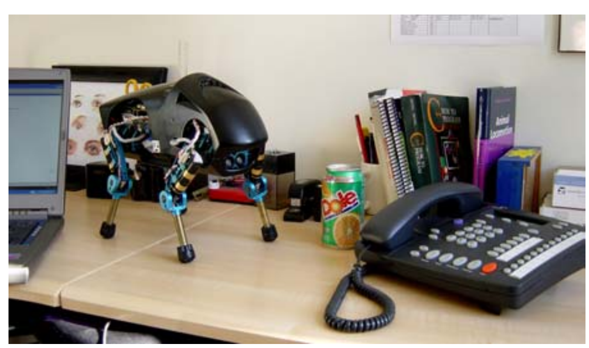

Figure 2. Little Dog Locomotion Platform

states the present parameter values for the Little Dog platform. The parameters shall be considered tentative and subject to change.

Table 2. Little Dog Parameters

Parameter	Value	Remarks
Number of Actuators	12	
Leg Length	6.3 in	
Body Length	10.6 in	
Body Width	3.5 in	
Body Height	10.2 in	
Belly Clearance	5.1 in	
Weight	6.6 lb	
Max Speed	10 in/sec	Ideal terrain
Leg Lengths per second	1.6	Ideal terrain
Body Lengths per second	0.9	Ideal terrain

Little Dog is a four-legged walking robot. It is controlled by an onboard computer and a Robot Control Station. The onboard computer performs all real-time control functions, while the Robot Control Station performs navigation, learning, and user interaction.

The Little Dog onboard computer performs the following functions:

- reads, calibrates and filters the robot sensors
- writes to actuators
- servos individual joints
- servos end-effector trajectories
- controls overall locomotion of the robot
- communicates with Robot Control Station
- logs data locally or on Robot Control Station via wireless communications

The onboard software is divided into three layers. The 'servo layer' samples and filters data for each sensor, performs servo functions for each motor, and determines the output to each motor. The servo layer runs at the highest priority and preempts all other activity on the Little Dog onboard computer, typically running once every 1 millisecond (1 kHz).

The 'communication layer' supports communications between Little Dog and the Robot Control Station. Communications are used for real-time control of the robot, real-time logging of data, user interaction, and to download controller modules to the control layer. Communications occur over 802.11b wireless Ethernet.

The 'control layer' controls Little Dog's behavior. It interprets data from the Little Dog's internal and navigational sensors, takes input from the Robot Control Station, applies one or more locomotion control algorithms, and specifies setpoints for the joint servos. We anticipate the control layer running every 20 to 50 milliseconds (50 to 20 Hz).

3.2 APPLICATION PROGRAMMER INTERFACE

The Robot Control Station communicates with Little Dog via the Little Dog C++ Application Programmer Interface (API).

The Little Dog API has the following functions:

- exchange data with the servo and control layers
- specify gains and setpoints to the joint servos
- read sensors
- write torques to actuators
- read and write internal state to robot
- specify data to be logged
- stream logged data from Little Dog and save to files

3.3 ROBOT CONTROL STATION

Each participating team will receive a Robot Control Station that has a wireless link to the Little Dog locomotion platform, an interface to the Terrain Board Cameras (see below), and two embedded high-performance Linux processors. One of the processors will be reserved for calculating the position of the Little Dog locomotion platform on the Terrain Board. The code running on this first processor will be fixed and will not be adjustable by the performers.

The other processor in the Robot Control Station will execute code to control the Little Dog locomotion platform, and to execute learning algorithms. The code on this second processor will be written by performers.

4 TERRAIN BOARDS

Each participating team will receive a Terrain Board as Government Furnished Equipment (GFE). A Terrain Board is a board with terrain features built onto it, similar to a scale-model railroad layout. The board serves as a surrogate for nature, allowing indoor operating unconstrained by weather, and provides identical challenges to the different performer teams.

Terrain Board design has not yet been completed. Preliminary designs call for two sections, each 4 feet wide and 8 feet long, to be combined to form a single 4-foot by 16-foot scale model. To meet the Phase I speed metric of 0.1 L/sec, Little Dog would need 320 sec to travel from one end to the other. (Again, L represents Little Dog leg length.) To meet the Phase II speed metric of 0.6 L/sec, 53 sec would be required for the end-to-end traverse.

To embody extreme natural terrain with large, irregular obstacles, Terrain Boards will include obstacles of varying size and with varying spacing. Representing the worst case for roughness, the tallest obstacle is expected to have height 0.9 L (5.7 in) measured from peak to valley. Representing the worst case for irregularity, the spatial frequency of footholds is expected to be between 0.5/L and 2/L.

The Robot Control Station will be augmented by a 3D tracking system that will "geolocate" the Little Dog within the Terrain Board. When combined with a terrain data file of the Terrain Board geometry (to be provided), this system will measure the 3D position and pose of the Little Dog relative to the Terrain Board coordinates, in near real time. It is expected that both the Terrain Board Camera system and the terrain data file will have sub-centimeter accuracy.

5 TEST AND EVALUATION

Progress in Phases I and II will be measured quantitatively through a series of official Government tests conducted at the Learning Locomotion Test Facility (LLTF). A DARPA-designated team independent of the developer teams will conduct the Government tests using a vehicle functionally identical to the GFE vehicles and a Robot Control Station essentially identical to the ones supplied to performers. Terrain boards used in the tests will not be physically identical to the GFE boards: the test boards will have similarly sized obstacles with some features similar to the GFE boards and some rather different. Performers are expected to have sufficiently adaptive algorithms that their systems will learn "on-the-fly" how to traverse these new obstacle types.

Developers will send their developed control software in the form of object code to the LLTF. There, operators will load the software onto a Robot Control Station, and command the associated Little Dog Platform to travel from one end of the board to the other. The operators will measure and document the performance of the system on multiple runs.

Government tests will measure the ability of the performer systems to learn from experience. These tests will take place about once a month, starting three months into the period of performance. The terrain encountered in a particular test is likely to have multiple obstacles with similar characteristics. For example, the test terrain board might include jumbles of flat surfaces that might appear in urban rubble, or large, rounded boulders that might occur in glacial deposits. During the course of the competition performer systems should learn how to traverse the particular class of obstacle types that are typical of that test. Thus, it is expected that performance will improve from one run to the next as the performer systems become familiar with the terrain and obstacles on the course. The systems will be able to store information gathered during a run using non-volatile storage on the Robot Control Station, which Government test operators will record. This information will be available to the performing team so that later runs may profit from the experience gained in earlier runs.

While it is expected performers will initially utilize locomotion that is statically stable, i.e. at least three legs on the ground and the center of mass within the triangle formed by the legs, performers are expected to also consider locomotion that is only dynamically stable so that balance is aided by the motion of the vehicle. Dynamically stable locomotion should allow big increases in vehicle speed, for example, by allowing the vehicle to "run" through relatively smooth stretches of the terrain board, or to use the vehicle's kinetic energy to "bound" from the top of one obstacle to the next.

6 PROGRAM SCOPE

6.1 LEARNING METHODOLOGIES

Performers will develop methods for training their vehicles to traverse large, irregular obstacles. Performers are encouraged to use whatever learning methodologies they believe will be most effective.

One possible approach to learning locomotion is illustrated in the following discussion:

In order to traverse an obstacle, the Robotic Control System needs to provide an appropriate sequence of actuator commands tailored to the initial position and pose of the vehicle. Symbolically, this means that the control system has to act like a function *L* that maps points in the input space, **P**, of pose, direction of a goal, and terrain, into points in the output space **S** of sequences of actuator commands:

$$\boldsymbol{s} = L(\boldsymbol{p}),\ \textit{where}\ \boldsymbol{s} \in \boldsymbol{S}\ \textit{and}\ \boldsymbol{p} \in \boldsymbol{P}$$

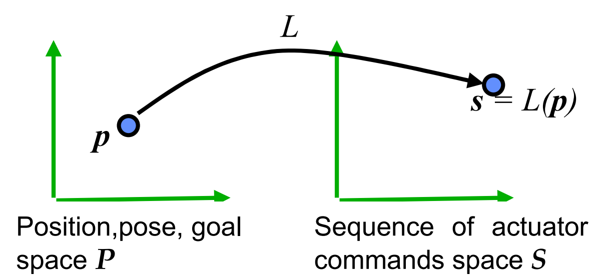

Figure 3. The Robotic Control System maps points in position-pose-terrain space ***P*** to points in sequence of actuator commands space ***S***.

It is unlikely that *L* has a simple analytic form, so the control system will have to have some way to empirically derive the appropriate values of *L* for the specific points in *P* that are of interest.

One way to determine values of $L(\boldsymbol{p})$ is to use the method of Reinforcement Learning. In a DARPA sponsored study, Reinforcement Learning was used to "teach" a simulated Little Dog to climb a high step. Learning to traverse a particular obstacle from a specific starting point (essentially one evaluation of *L*) took several minutes to compute on a workstation.

This Reinforcement Learning code was not optimized and the researchers estimate that with faster executing code, evaluations of *L* for perhaps 1000 different values of $\boldsymbol{p}$ could be done in a single day. However, even at this execution speed, forward motion of the vehicle would be excruciatingly slow.

The good news is that *L* should be piecewise continuous, so that if $\boldsymbol{s} = L(\boldsymbol{p})$ is known, then $L(\boldsymbol{p}+\varepsilon)$ is expected to be close to $L(\boldsymbol{p})$ for small ε. This means that if the control system has a sufficiently dense sampling of empirically derived values of *L* for some region in ***P***, the control system should be able to interpolate and get pretty good approximations for *L* over that entire region in ***P***. See Figure 4.

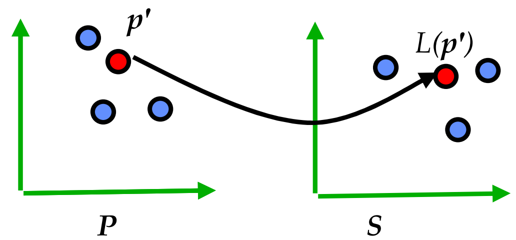

Figure 4. If *L(**p**)* has already been computed for a number of points in ***P*** (blue dots), then *L(**p'**)* can be estimated through interpolation (red dot).

When the vehicle confronts an obstacle, corresponding to being at point ***p'***, it is unlikely that the control system has previously, computed *L(**p'**)*, but if it has already computed L for points near ***p'*** then it can interpolate and get a good estimation of *L(**p'**)* that can be refined by doing a small (presumably fast) amount of reinforcement learning starting from the interpolated estimate.

The above strategy requires that the control system has a good way of estimating *L(**p'**)*. One way to do this is to use the method of learning from example: The previously computed values of L for various values of p can be used to form a training set of [***p*** , *L(**p**)*] pairs. These pairs could be then used for training a neural net, a support vector machine, or a nearest neighbors estimator.

This discussion sketches a plausible approach to using learning methods for locomotion. Performers are free to pursue (or not pursue) such an approach. It has been included in this PIP simply to inspire potential proposers.

6.2 INNOVATION

Proposed research should investigate innovative learning approaches and techniques that lead to, enable revolutionary advances in the state of the art. Proposals should be for research that substantially contributes to the goal of Learned Locomotion over extreme terrain. Specifically excluded is research that primarily results in minor evolutionary improvement to the existing state of practice or focuses on special-purpose systems, narrow applications, and/or handcrafted coding.

6.3 MEETINGS

In addition to participating in all scheduled competitions, performers must participate in any Principal Investigator meetings held twice a year as part of the Learning Locomotion program. At these meetings, it is expected that each performer will report on the design and implementation of the learning methods. It is anticipated that these reports will include analysis of why the system made key decisions during test runs, analysis based in part upon data collected and logged during the test runs.

6.4 COOPERATION

Cooperation between performers is encouraged. Cooperation can include, but is not restricted to, sharing learned behaviors, as well as learning algorithms. Prior to tests, performers must disclose any software they have incorporated from other performers into their systems as well as any training data that they received.

6.5 OBJECT CODE

DARPA intends to provide a repository for the object code submitted to the LLTF prior to each competition. This object code may be shared among performers.

7 GENERAL INFORMATION

Proposals not meeting the format described in this pamphlet may not be reviewed. Proposals MUST NOT be submitted by fax or e-mail; any so sent will be disregarded. This notice, in conjunction with the BAA BAA 05-25 FedBizOpps Announcement and all references, constitutes the total BAA.

A Frequently Asked Questions (FAQ) list may be provided. The URL for the FAQ will be specified on the DARPA/IPTO Solicitation page.

An Industry Day will be held for the Learning Locomotion program. Details and registration information can be found at: http://safe.sysplan.com/locomotion

No additional information is available, nor will a formal Request for Proposal (RFP) or other solicitation regarding this announcement be issued. Requests for same will be disregarded.

All responsible sources capable of satisfying the Government's needs may submit a proposal that shall be considered by DARPA. Historically Black Colleges and Universities (HBCUs) and Minority Institutions (MIs) are encouraged to submit proposals and join others in submitting proposals. However, no portion of this BAA will be set aside for HBCU and MI participation due to the impracticality of reserving discrete or severable areas of this research for exclusive competition among these entities.

The Government anticipates that proposals submitted under this BAA will be unclassified. In the event that a proposer chooses to submit a classified proposal or submit any documentation that may be classified, the following information is applicable. Security classification guidance on a DD Form 254 will not be provided at this time since DARPA is soliciting ideas only. After reviewing incoming proposals, if a determination is made that the award instrument may result in access to classified information, a DD Form 254 will be issued and attached as part of the award. Proposers choosing to submit a classified proposal must first receive permission from the Original Classification Authority to use their information in replying to this BAA. Applicable classification guide(s) should be submitted to ensure that the proposal is protected appropriately.

8 SUBMISSION PROCESS

This BAA requires completion of an online Cover Sheet for each Proposal prior to submission. To do so, the offeror must go to **http://www.dyncorp-**

is.com/BAA/index.asp?BAAid=05-25 and follow the instructions there. Each offeror is responsible for printing the BAA Confirmation Sheet and attaching it to every copy. The Confirmation Sheet should be the first page of the Proposal. If an offeror intends to submit more than one Proposal, a unique UserId and password must be used in creating each Cover Sheet. Failure to comply with these submission procedures may result in the submission not being evaluated.

Proposers must submit the original and 4 copies (5 total) of the full proposal and 2 electronic copies (i.e., 2 separate disks) of the full proposal (in PDF or Microsoft Word 2000 for IBM-compatible format on a 3.5-inch floppy disk, 100 MB Iomega Zip disk or CD). Mac-formatted disks will not be accepted. Each disk must be clearly labeled with BAA BAA 05-25, proposer organization, proposal title (short title recommended) and "Copy ___ of 2". The full proposal (original and designated number of hard and electronic copies) must be submitted in time to reach DARPA by 12:00 PM (ET) April 15, 2005, in order to be considered during the initial evaluation phase. However, BAA 05-25, Learning Locomotion will remain open until 12:00 NOON (ET) March 1, 2006. Thus, proposals may be submitted at any time from issuance of this BAA through March 1, 2006. While the proposals submitted after the April 15, 2005 deadline will be evaluated by the Government, proposers should keep in mind that the likelihood of funding such proposals is less than for those proposals submitted in connection with the initial evaluation and award schedule. DARPA will acknowledge receipt of submissions and assign control numbers that should be used in all further correspondence regarding proposals.

Restrictive notices notwithstanding, proposals may be handled for administrative purposes by support contractors. These support contractors are prohibited from competition in DARPA technical research and are bound by appropriate non-disclosure requirements. Input on technical aspects of the proposals may be solicited by DARPA from non-Government consultants /experts who are also bound by appropriate non-disclosure requirements. However, non-Government technical consultants/experts will not have access to proposals that are labeled by their offerors as “Government Only.” Use of non-Government personnel is covered in FAR 37.203(d).

9 REPORTING REQUIREMENTS AND PROCEDURES

The Award Document for each proposal selected and funded will contain a mandatory requirement for submission of DARPA/IPTO Quarterly Status Reports and an Annual Project Summary Report. These reports will be electronically submitted by each awardee under this BAA via the DARPA/IPTO Technical - Financial Information Management System (T-FIMS). The T-FIMS URL will be furnished by the Government upon award. Detailed data requirements can be found in the Data Item Description (DID) DI-MISC-81612A available on the Government's ASSIST database (http://www.darpa.mil/leaving.asp?url=http://assist.daps.dla.mil/quicksearch/).

10 PROPOSAL FORMAT

Proposals shall consist of a cover page, a technical volume, and a cost volume. The submission of other supporting materials—including bibliographies, technical papers, and research notes—along with the proposal is strongly discouraged.

A "page" is 8-1/2 by 11 inches with type not smaller than 12 point, and with text on one side only.

10.1 COVER PAGE

The cover page shall be the confirmation sheet referenced in Section 8 Submission Process.

10.2 VOLUME I. TECHNICAL

This volume provides the detailed discussion of the proposed work necessary to enable an in-depth review of the specific technical and management issues. Specific attention must be given to addressing both the risk and payoff of the proposed work that make it desirable to DARPA.

The Technical Volume shall not exceed 17 pages, and shall include sections A through H, each beginning on a new page. Maximum page lengths for each section are shown in braces { } below, where applicable.

A. Innovative claims for the proposed research {1 Page}. This page is the centerpiece of the proposal and should succinctly describe the unique proposed contribution.

B. Proposal Roadmap {1 Page}. The roadmap provides a top-level view of the content and structure of the proposal. It contains a synopsis (or "sound bite") for each of the four areas defined below. It is important to make the synopses as explicit and informative as possible. The roadmap must also cross-reference the proposal page number(s) where each area is elaborated. The four roadmap areas are:

1. Critical technical barriers (i.e., technical limitations that have, in the past, prevented autonomous locomotion over extreme terrain).
2. Main elements of the proposed approach.
3. The proposer should clearly show the rationale that their approach will overcome the technical barriers. ("We have a good team and good technology" is not a useful statement.)
4. Cost of the proposed effort for each performance year.

C. Statement of Work {3 Pages}. Detailed statement of work, written in plain English, outlining the scope of the effort and citing specific tasks to be performed, references to specific subcontractors if applicable, and specific contractor requirements.

D. Technical Approach:

1. Detailed Description of Technical Approach {7 Pages}. Provide detailed description of technical approach that will be used in this project to achieve

learned locomotion over extreme terrain. This section should address in depth issues such as (but not necessarily limited to) learning time, efficient use of previously learned behaviors, closing the control loops using data from the Terrain Board Cameras, and how to take advantage of vehicle dynamics.

2. Comparison with Current Technology {1 Page}. Describe state-of-the-art approaches and the limitations within the context of the problem area addressed by this research.

E. Schedule: {1 Page}. Provide a graphic representation of project schedule including detail down to the individual effort level. This should include, but not be limited to, a multi-phase development plan, which demonstrates a clear understanding of the proposed research; and a plan for periodic and increasingly robust experiments over the project life that will show applicability to the overall program concept.

F. Deliverables Description {1 Page}. List and provide detailed description for each proposed deliverable. Include in this section all proprietary claims to results, prototypes, or systems supporting and/or necessary for the use of the research, results, and/or prototype. If there are no proprietary claims, this should be stated. The offeror must submit a separate list of all technical data or computer software that will be furnished to the Government with other than unlimited rights (see DFARS 227.) Specify receiving organization and expected delivery date for each deliverable.

G. Personnel and Qualifications {1 Page}. List of key personnel, concise summary of their qualifications, and discussion of previous accomplishments and work in robotics, machine learning, or closely related research areas. Indicate the level of effort to be expended by each person during each contract year and other (current and proposed) major sources of support for them and/or commitments of their efforts. DARPA expects all key personnel associated with a proposal to make substantial time commitment to the proposed activity.

H. Facilities {1 Page}. Description of the facilities that would be used for the proposed effort.

10.3 VOLUME II. COST

Cost proposals are subject to no page limits, and shall provide a detailed cost breakdown of all direct costs, including cost by task, with breakdown into accounting categories (labor, material, travel, computer, subcontracting costs, labor and overhead rates, and equipment), for the entire contract and for each calendar year, divided into quarters. Participation in Phase II should be indicated as a contract option with a separate cost estimate. The Government anticipates multiple awards of approximately $600K-$800K per Phase. Continued funding into Phase II will be based on achieving the Phase I performance criteria.

Show your costs in a Milestone Chart against the listed project milestones. Use Months after Contract (MAC) for designations for all dates. Proposed Milestones:

Phase I	MS	Event	Requirement	End Date
	1	Kickoff Meeting	Participate in Kickoff Take delivery of vehicle	2 MAC
	2	Test 1 – Phase I	Upload code for Test 1	5 MAC
	3	PI Meeting	Participate in PI Meeting Upload code for 4 of 5 tests	10 MAC
	4	End of Phase 1	Participate in PI Meeting Upload code for 4 of 5 tests Phase I Final Report	15 MAC
Phase II	5	Test 1 – Phase II	Upload code for Ph II Test 1	17 MAC
	6	Monthly Testing	Upload code for 3 of 4 tests	21 MAC
	7	PI Meeting	Participate in PI Meeting Upload code for 3 of 4 tests	25 MAC
	8	PI Meeting	Participate in PI Meeting Upload code for 3 of 4 tests	29 MAC
	9	End of Phase II	Participate in PI Meeting Upload code for 3 of 4 tests Phase I Final Report	33 MAC

Offerors should expect to attend a kickoff meeting and semi-annual Principal Investigator (PI) meetings to provide technical information to DARPA, and participate in other coordination meetings via teleconference or Video Teleconference (VTC). Funding to support these various efforts should be included in technology project bids

IMPORTANT NOTE: IF THE OFFEROR DOES NOT COMPLY WITH THE ABOVE STATED REQUIREMENTS, THE PROPOSAL WILL BE REJECTED.

10.4 ORGANIZATIONAL CONFLICT OF INTEREST

Awards made under this BAA may be subject to the provisions of the Federal Acquisition Regulation (FAR) Subpart 9.5, Organizational Conflict of Interest. All offerors and proposed subcontractors must affirmatively state whether they are supporting any DARPA technical office(s) through an active contract or subcontract. All affirmations must state which office(s) the offeror supports, and identify the prime contract number. Affirmations should be furnished at the time of proposal submission. All facts relevant to the existence or potential existence of organizational conflicts of interest, as that term is defined in FAR 2.101, must be disclosed in Volume I, Section G. of the proposal, organized by task and year. This disclosure shall include a description of the action the Contractor has taken, or proposes to take, to avoid, neutralize, or mitigate such conflict.

11 EVALUATION AND FUNDING PROCESSES

Proposals will not be evaluated against each other, since they are not submitted in accordance with a common work statement. DARPA's intent is to review proposals as soon as possible after they arrive; however, proposals may be reviewed periodically for administrative reasons. For evaluation purposes, a proposal is the document described in Proposal Format. Other supporting or background materials submitted with the proposal will be considered for the reviewer's convenience only and not considered as part of the proposal.

Evaluation of proposals will be accomplished through a scientific review of each proposal using the following criteria, which are listed in descending order of relative importance:

(1) Overall Scientific and Technical Merit: The overall scientific and technical merit must be clearly identifiable and compelling. The technical concepts should be clearly defined and developed. The technical approach must be sufficiently detailed to support the proposed concepts and technical claims. Evaluation will also consider the system integration and management plan.

(2) Control Issues: Demonstrated understanding of the issues in controlling a high degree-of-freedom quadruped robot

(3) Implement Methods: Demonstrated understanding of how to implement machine learning methods in real-world problems

(4) Offeror's Capabilities and Related Experience: The qualifications, capabilities, and demonstrated achievements of the proposed principals and other key personnel for the primary and subcontractor organizations must be clearly shown.

(5) Cost Realism: The overall estimated costs should be clearly justified and appropriate for the technical complexity of the effort. Evaluation will consider the value of the research to the Government and the extent to which the proposed management plan will allocate resources to achieve the capabilities proposed.

The Government reserves the right to select for award all, some, or none of the proposals received. Proposals identified for funding may result in a contract, grant, cooperative agreement, or other transaction depending upon the nature of the work proposed, the required degree of interaction between parties, and other factors. If warranted, portions of resulting awards may be segregated into pre-priced options.

12 ADMINISTRATIVE ADDRESSES

The administrative addresses for this BAA are the following:
Fax: 703-741-7804, addressed to: DARPA/IPTO, BAA 05-25
Electronic Mail: BAA 05-25 @darpa.mil
Electronic File Retrieval: http://www.darpa.mil/ipto/Solicitations/solicitations.htm

Mail: DARPA/IPTO
ATTN: BAA 05-25
3701 N. Fairfax Drive
Arlington, VA 22203-1714