

Adaptive Resource Modelling for Autonomous Planning and Scheduling

Andrew Carrel and Phil Palmer

Surrey Space Centre
University of Surrey
Guildford GU2 7XH
United Kingdom

Abstract

Autonomous planning and scheduling requires models of how spacecraft operations will affect resources. In practice the true behaviour may deviate from the model, either due to modelling inaccuracies or a change in behaviour. In this paper a neural network is used to learn about such modelling errors and provide corrections that can be used for forward planning. This enables adaptation to unexpected changes in the way resources are consumed, such as may occur when a subsystem's performance becomes degraded.

This learning method has been incorporated into the NEAT autonomous planner to enable it to adapt its resource models. Comparisons of NEAT's resource management performance with and without this new component have shown that using a neural network in this way greatly improves the robustness of autonomous operations when faults develop.

Introduction

The NEAT algorithm has previously been presented (Carrel & Palmer 2005) as an effective approach to onboard autonomy which could produce significant performance improvements over more conventional spacecraft operations. This Near-optimal, Evolutionary, Autonomous Task-manager (NEAT) optimises its planned schedule of operations using an evolutionary algorithm, whilst enforcing temporal and resource constraints. Operations and resources are defined using a generic format such that different types of operation can be scheduled and multiple resources can be managed. This will make it possible to apply NEAT to a wide range of applications.

NEAT works on an entirely iterative basis, with the operations schedule evolving towards the optimum continuously. When a decision is to be made the best solution in the current population is used. Since the evolutionary process is continuous, NEAT is able to adapt to spurious events such as a failed operation or an unexpected loss of resource. Previous results have shown this method to be effective at repeatedly re-optimising the operations schedule in the presence of such events (Carrel & Palmer 2005).

To be able to manage future resource levels, any planning algorithm must have models describing the expected resource use of each operation and the availability of these

resources at future times. In practice however, these models will generally be imperfect, either due to simplifying assumptions of inherent unpredictability. For this reason NEAT updates its estimates of resource levels as measurements become available and adapts its plans as necessary. However, this feedback does not learn about systematic modelling errors such as may occur when the spacecraft changes its behaviour. For example, if a solar array is degraded by micro-meteorite impacts, the available electrical power will be systematically less than the model may predict.

In this work the NEAT algorithm is augmented with a neural network which is able to learn about systematic changes in the availability or consumption of resources. Measured resource levels are compared with those predicted during planning to find systematic modelling errors (see Figure 1). The neural network is then trained to predict these errors such that corrections can be applied for future planning. The training data is continuously updated such that the network adapts with changing resource behaviour.

This system has been applied to the case of scheduling imaging operations on the UK-DMC Earth observation satellite. Faults were simulated to create changes in the way operations affect resource levels. The neural network was then required to adapt so that NEAT could continue to make effective plans autonomously.

Generic Resources and Activities

A spacecraft may have many different resources, some of which are tangible such as fuel, and some of which are related to its circumstances, such as the availability of a ground station. For it to be possible to create a generic resource management system it is first necessary to divide resources into well-defined categories.

Two resource categories are used here:

Depletable resources are those that remain depleted even after an activity has finished using them. Examples include fuel and electrical energy. In some cases, such as electrical energy, it may be possible for an activity to replenish rather than consume the resource. Only a finite quantity of this type of resource will be available at any one time and this quantity will never exceed a fixed upper limit, for example a vehicle's battery capacity in the

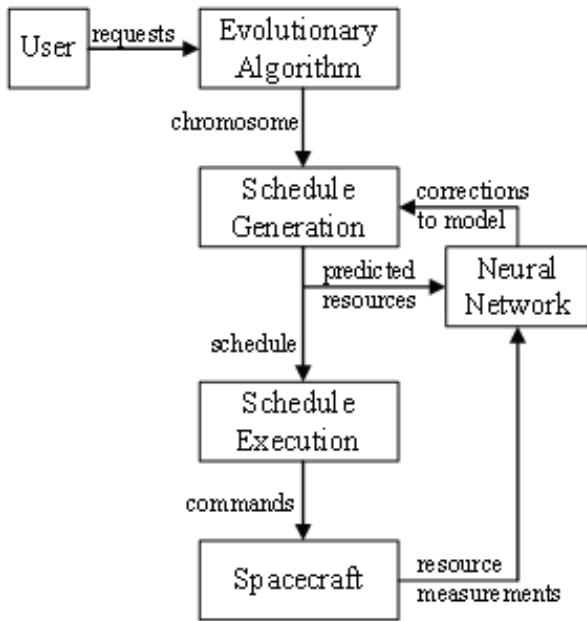


Figure 1: Scheme for adapting to changing resource behaviour. Measured resource levels are compared with predicted levels to create resource model corrections.

case of electrical energy. Depletable resources are depleted and replenished at finite rates over finite time intervals such that there are no step changes in the available quantity of the resource.

Non-depletable resources are only depleted for the period in which they are in use, after which they return to their original level. An example is computer processing power. Some activities may make more of this type of resource available for their duration rather than using the resource. The availability of this type of resource must not become negative and has fixed maximum limit. The effect of an activity using this type of resource will be a step change in the resource level at the start of the activity, which is reversed when the activity is complete.

These definitions are similar to those used in the the ASPEN system (Fukunaga *et al.* 1997; Chien *et al.* 2000). In that work the two are differentiated by the fact that a depletable resource “remains used even after the activity” (Chien *et al.* 2000), whilst “a non-depletable resource is used only for the duration of the activity”. In a separate work, Brambilla *et al.* (Brambilla *et al.* 2005) also refer to these two resource types and define them in much the same way. Similarly, Funase & Nakasuka (Funase & Nakasuka 2005) consider “consumable” and “reusable” resources. Lemai & Ingrand (Lemai & Ingrand 2004) do not categorise resources but rather allow actions to “use, consume and produce” resources. When a resource is ‘used’ it is borrowed over an interval, producing a similar behaviour to a non-depletable resource. Consumption of a resource has an effect after the action is complete, similar to the effect

generated by using a depletable resource.

Whilst the definitions of depletable and non-depletable resources given here are broadly the same as in other works, one difference is worth noting. Elsewhere it is assumed for planning purposes that an activity’s depletable resource requirements are completely used from the instant the activity begins. In this work, however, activities consume (or produce) depletable resources at a defined rate over their duration. This approach is particularly useful when a depletable resource is being consumed or generated at a low rate over a relatively long period since the expected availability of the resource will be much more accurate.




Figure 2: Modelling of non-linear charging of a resource using three linear activities.

Where an activity is expected to consume or produce a depletable resource in a non-linear fashion, multiple activities can be used to better represent the resource profile. For example the rate of charge (power) generated by a solar array will usually be a sine wave, which is clearly non-linear. The charge level over this period can, however, be modelled reasonably accurately using a few activities, as shown in Figure 2.

Both depletable and non-depletable resources have minimum and maximum levels. It is therefore straightforward to rescale each resource such that its minimum and maximum levels are zero and one respectively. This is very useful when dealing with generic resources since now the current state of any resource can be represented by a value in the range zero to one.

An activity is considered to be any action that partly or completely achieves a goal, or otherwise affects the spacecraft. An operation is then simply the group of all the activities that relate to a goal the spacecraft has been set. Operations are mutually exclusive sets of activities. When all the constituent activities of such a set are successfully completed then the relevant goal will have been achieved.

To create an operation request (i.e. set the spacecraft a goal) it is first necessary to list all the activities that will be required. With each of these a finite set of one or more time windows must be given that specify when the activity may

Constraint	Description
$S_A > S_B$	Start activity A after B starts
$S_A > F_B$	Start activity A after B finishes
$F_A > S_B$	Finish activity A after B starts
$F_A > F_B$	Finish activity A after B finishes
$S_A < S_B$	Start activity A before B starts
$S_A < F_B$	Start activity A before B finishes
$F_A < S_B$	Finish activity A before B starts
$F_A < F_B$	Finish activity A before B finishes
$S_A = S_B$	Start activity A as B starts
$S_A = F_B$	Start activity A as B finishes
$F_A = S_B$	Finish activity A as B starts
$F_A = F_B$	Finish activity A as B finishes

Table 1: Possible relative temporal constraints between the start or finish time of activity A (S_A, F_A) and the start or finish time of activity B (S_B, F_B).

be scheduled. Activities may also have temporal constraints that are relative to other activities within the same operation (see Table 1). For example image data must be transmitted to a ground station after the image has been captured and stored. If an activity's duration is not determined by relative constraints then it is specified explicitly.

For each activity the required amount of each of the spacecraft's resources (or the amount generated) is specified. For depletable resources this will be the rate at which the resource will be consumed (or produced). Some activities may have to be included to account for essential or unavoidable changes in resource availability. For example charging of batteries using solar arrays is generally essential and some power consumption is usually unavoidable. These activities may not relate directly to high-level goals and so are goals in themselves. Operation requests must be created for individual or groups of activities of this kind so that forward planning is done in light of all expected resource changes.

Each goal is given a priority based on its perceived importance. Firstly goals are split into two groups, essential and non-essential. Essential goals relate to those activities just described that either must be performed to maintain the spacecraft's health or have unavoidable impacts on resources. All of the operations relating to these goals will be scheduled with equal priority. This should be possible without conflict, otherwise there is something wrong with the spacecraft. All other, non-essential goals are given a user defined priority, which is taken into account when optimising the schedule.

Schedule Generation and Optimisation

The continuous evolutionary optimisation of the operations schedule enables the spacecraft to be highly adaptive, however it is important to ensure that all the temporal and resource constraints are satisfied. Also the repeated projection of different planning options requires that schedules can be generated quickly. All the essential operations are used to create a baseline schedule. The optimisation component then searches the different orders in which the non-essential

operations may be considered for placement in the schedule. Each operation consists of a set of activities, which are considered for placement in the schedule in a fixed order. These orderings do *not* mean that activities must be scheduled in a set time order.

An operations schedule is represented by a set of time-ordered events, representing the start and end of each activity. Resource availability and depletable resource rates are stored with each event. Depletable resource levels can then be projected directly from one event to the next without fixed time-step numerical integration to create a profile of the expected future resource availability (see Figure 3).




Figure 3: Projection of one depletable resource (D_1) and two non-depletable resources (N_1 & N_2) using schedule events.

The process of placing an operation in the schedule is based on backtracking. This is designed to ensure that all options have been exhausted before an operation is rejected. This does not, however, go as far as considering moving any activities relating to operations that have already been placed in the schedule, since these options are dealt with by the evolutionary optimisation.

Scheduling an operation simply requires finding a place in the schedule for each of its constituent activities, $a = 1 \dots N_a$, such that all temporal and resource constraints are satisfied. The method for working through the activities is shown in Figure 4. The process starts with the first activity ($a = 1$), applying the function $S(a)$, which attempts to place it in the schedule. If this is successful then the algorithm moves on to the next activity and repeats this scheduling until either an activity cannot be placed or all the activities are scheduled and so the process is complete. This is the forward loop marked 'F' in the Figure.

If the function $S(a)$ fails to place an activity then the algorithm moves back to the previous activity, unless this is already the first activity ($a = 1$). In the latter case, the process terminates having failed either because it was not ever possible to place the first activity or after an exhaustive search in which many possible times were considered but in each case one of the other activities could not be sched-

Neural Network Learning

The neural network used here consists of an input layer of size n_r (the total number of resources), a layer of $n_h = 2n_r$ hidden neurons and a layer of n_r output neurons as shown in Figure 6. The network has one input and one output for each resource being managed. For non-depletable resources the inputs are the expected normalised resource levels whereas for depletable resources the nominal rate of consumption or production is used. The outputs of the network are corrections to be applied to the non-depletable resource levels and the depletable resource rates.

In the case study considered in the next section it is corrections to the rate of consumption or generation of the depletable resources that is of particular interest, since these resources are the most difficult to manage. The non-depletable resource levels give an indication of which subsystems are in use. When trained, the neural network will therefore be able to determine which subsystems, or combinations of subsystems, are in use when deviant resource behaviour occurs. This enables the network to predict such behaviour in future and calculate the necessary model corrections. Other inputs to the network could be incorporated, such as the types of activity being undertaken at the time.




Figure 6: Neural network structure for two resources.

Both the inputs and outputs are rescaled to ensure that the network only has to deal with values comfortably in the range zero to one. For non-depletable resources input values originally in the full zero to one range are rescaled to the range 0.1 to 0.9. Depletable resource rates are not constrained to a particular range and so each raw depletable input D is rescaled to x_D using:

$$x_D = \frac{1}{1 + e^{-\sigma D}} \quad (1)$$

The scale factor σ is set to a value appropriate to the problem in question. These scalings are applied in reverse to the network outputs to get the final correction values.

Calculation of Weights

The hidden neurons are fully connected to the input layer with weight matrix W_1 and the output neurons are fully connected to the outputs \vec{u}_1 of the hidden neurons with weight matrix W_2 . The activation vectors \vec{a}_1 and \vec{a}_2 of the hidden and output layers respectively are given by:

$$\vec{a}_1 = W_1 \vec{x} \quad (2)$$

$$\vec{a}_2 = W_2 \vec{u}_1 \quad (3)$$

For each neuron the output is related to the activation by:

$$u = \frac{1}{1 + e^{-a}} \quad (4)$$

This is used to generate the output vectors \vec{u}_1 and \vec{u}_2 from the activation vectors \vec{a}_1 and \vec{a}_2 .

The network is trained using backpropagation (Rumelhart & et al 1986) to adjust the weights, which are initially set to zero. For the i th set of training inputs \vec{x}_i and outputs \vec{t}_i , the incremental adjustments to the weights are given by:

$$\delta W_{1i} = \beta \vec{d}_1 \vec{x}_i^T \quad (5)$$

$$\delta W_{2i} = \beta \vec{d}_2 \vec{u}_1^T \quad (6)$$

where:

$$\vec{d}_1 = W_2^T \vec{d}_2 \otimes \vec{u}_1 \otimes (\vec{1} - \vec{u}_1) \quad (7)$$

$$\vec{d}_2 = (\vec{t}_i - \vec{u}_2) \otimes \vec{u}_2 \otimes (\vec{1} - \vec{u}_2) \quad (8)$$

Here $\vec{1}$ refers to a vector where each element equals 1 and the symbol \otimes refers to element by element multiplication of two vectors to get a third vector of the same length. β is the learning rate. The weight adjustments are summed over the whole set of training data:

$$\delta W_1 = \sum_{i=1}^{N_{train}} \delta W_{1i} \quad (9)$$

$$\delta W_2 = \sum_{i=1}^{N_{train}} \delta W_{2i} \quad (10)$$

These adjustment matrices are then added to the weight matrices. This process is repeated until the network outputs are a close match to the test data set of N_{test} elements or until these weight changes have been calculated and applied for a maximum C_{train} cycles. The learning rate β starts small with $\beta = \beta_0 / C_{train}$ and increases linearly up to a maximum value β_0 , which is only reached if the maximum C_{train} cycles are employed.

The network outputs and the test data are considered to be a close match when the RMS error between them is below a threshold ϵ_{max} . The error is given by:

$$\epsilon = \sqrt{\frac{1}{n_r N_{test}} \sum_{i=1}^{N_{test}} \sum_{r=1}^{n_r} (t_{i,r} - u_{2,r})^2} \quad (11)$$

where \vec{t}_i is the output vector for the i th element of the test data set and \vec{u}_2 is the vector of outputs produced by the network when the corresponding test data inputs are applied.

Operation	Non-depletable Resources					Depletable Resource Rates		
	Camera	Antenna	Gnd Station	Solar Arrays	Power	Bat. SOC	SSDR0	SSDR1
Solar Power Gen.	0.0	0.0	0.0	-1.0	0.0	+57.3W	0.0	0.0
Platform maint'nce	0.0	0.0	0.0	0.0	0.0	-23.1W	0.0	0.0
GPS orbit det.	0.0	0.0	0.0	0.0	1.0	-5.7W	0.0	0.0
G.S. visibility	0.0	0.0	+1.0	0.0	0.0	0.0W	0.0	0.0
Image capture	-1.0	0.0	0.0	0.0	1.0	-32.0W	≤ 0.0	≤ 0.0
Downlink	0.0	-1.0	-1.0	0.0	1.0	-39.1W	+3.9Mbps	+3.9Mbps

Table 2: Resource requirements for UK-DMC operations. Positive values indicate generation of a resource, whilst negative values indicate consumption. The rate of memory consumption during image capture depends of the geometry with respect to the target.

Resource Model Adaptation

Once the neural network has been trained it can be used to modify the expected resource use of activities during the planning process. Each event in the schedule stores the nominal availability of non-depletable resources and rates of consumption or production of depletable resources. These values are passed as the input vector \vec{x} to the neural network such that the output vector \vec{u}_2 can then be used to generate corrected values, which are stored alongside the nominal values at each event. In this way the corrected values can be used for projecting resource levels for planning and the nominal values can still be used for comparison with measurements to create training data.

Training the Network

Training data is accumulated each time an event in the schedule is reached. At these points the true resource levels are measured just before the scheduled event is acted upon. For non-depletable resources the difference between the expected and measured levels is recorded. For depletable resources the rate at which the resource was actually used is calculated from the measurement at the last event, the current measurement and the elapsed time. This is compared to the nominal rate expected over that interval and the necessary correction is stored. The corresponding network inputs that are stored are the nominal non-depletable resource levels and depletable resource rates expected over the interval between the two events.

Each time a new piece of training data is added to the data set, the oldest data element is deleted and the network weights are readjusted. In this way the network is always trained using a recent set of data of a fixed size. This means that the resource use model used for forward planning is continually adapting in light of recent resource measurements. The total data set is of size N_{data} , with half of the data (N_{train} elements) used for training and the remaining N_{test} elements used for testing the network.

The size of the data set will determine how responsive the neural network is to changes in resource behaviour. If the data set is very large then it may take a long time for new measurements to supplant old data. However, if the training data set is too small then the neural network may learn about noise or short term fluctuations and produce spurious corrections.

Test Case: UK-DMC Imaging

SSTL's UK-DMC satellite is a 700km altitude Sun-synchronous Earth observation satellite operated from SSTL in Guildford, UK. It is one of five satellites that form the Disaster Monitoring Constellation (Chu *et al.* 2000), which provides medium resolution imagery with global coverage and a daily revisit period.

Here NEAT is applied to the UK-DMC case both with and without the neural network learning. Three versions of NEAT are tested: open-loop planning, planning with feedback of measured resource levels and planning with the neural network learning described above.

UK-DMC has three depletable resources that need to be managed: two solid state data recorders (SSDR0 and SSDR1) and the battery state-of-charge (SOC). The capacities of SSDR0 and SSDR1 are 1024 Mbytes and 512 Mbytes respectively. The battery consists of 22 4000mAh NiCd cells in series and has a nominal voltage of 28V. Therefore the capacity is taken to be 403.2 kJ, however its state of charge should not fall below 85% to prevent degradation of performance. A charging efficiency of 92% is used to account for losses in the power system.

Five non-depletable resources are also included, all of which are 'atomic' here in that they are either completely in use or fully available. These are used to represent the camera, antenna, solar arrays, the availability of the ground station and to limit power use. The last of these allows only one GPS, image capture or downlink operation to occur at any one time.

The different operations to be scheduled are listed in Table 2 along with the impact they have on each resource. In this case each operation consists of just one activity. Four essential operations represent the routine functions onboard the satellite. Solar power generation is active whenever the spacecraft is in sunlight and the core platform operation encompasses the continuous maintenance activities such as attitude control. 10 minute GPS orbit determination operations are scheduled at 110 minute intervals such that GPS measurements are progressively staggered around the orbit (this is in line with current operations procedures). The non-depletable resource representing ground station availability is initialised as unavailable (i.e. set to zero) and operations are scheduled that provide availability for each visibility window.

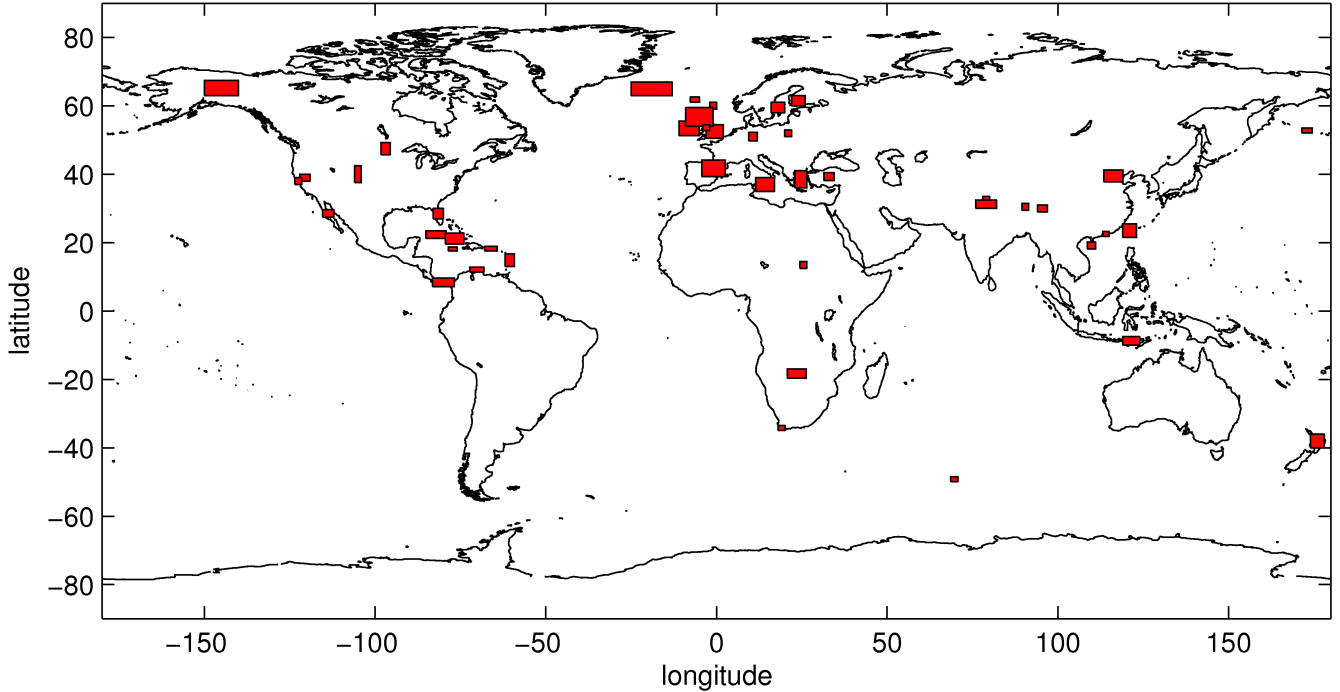


Figure 7: Target image areas. Some of these targets occur more than once in the set of targets to be imaged.

Image capture and downlink operations are the two types of non-essential operations. The evolutionary algorithm searches for the optimal way in which to schedule these operations. User requests take the form of a target latitude and longitude, the latitude and longitude extent of the area to be imaged and a time by which the image is to be taken. ImPredict (Mai & Palmer 2001; Wu, Brewer, & Palmer 2001) is used to calculate time windows for image capture and for ground station visibility. This is a fast overpass calculation algorithm, based on the analytical epicycle orbit model (Hashisa & Palmer 2001), developed at Surrey. The memory rates are variable due to the nature of the UK-DMC imaging system. The full width of the field of view is around 51.9° , however this is split between two CCD arrays, one port and one starboard facing. These feed directly into SSDR0 and SSDR1 respectively. Therefore since UK-DMC is a nadir-pointing satellite the proportions of the image that will be stored by the two data recorders will vary from one imaging opportunity to the next.

To provide a realistic test scenario a set of imaging targets has been taken from the imaging history of UK-DMC. The 94 targets imaged by UK-DMC during the month of September 2004 are used as the image request targets for testing NEAT in this case. These target areas are spread widely over the globe and are shown in Fig. 7. The scheduling for this period was performed by ground station staff using the existing Mission Planning System, which also calculates imaging windows using ImPredict. The Mission Planning System automatically enforces operational constraints such as resource limits but maximising the throughput of the sys-

tem is conducted manually.

The same 94 image requests were sent to NEAT in a simulation of UK-DMC over the same period, with the orbit set to its state at that time. Requests were sent such that there were always 50 outstanding and should all 94 requests have been sent the simulation started sending the requests again from the top of the list but with a much lower priority ($P = 1$ as opposed to $P = 5$).

For this application the fitness, Q , of a trial schedule in the evolving population is given by:

$$Q = \sum_{i=1}^L P_i e^{-\gamma \tau_i / T} \quad (12)$$

where P_i is the priority of the i th image request, τ_i is the time taken to return the i th image, T is the maximum image return time that will be accepted and L is the chromosome length, corresponding to the number of outstanding requests. This function ensures that higher priority operations are given more weight. Also the parameter γ determines how much pressure there is to select earlier imaging times where possible ($\gamma = 4.0$ is used).

When NEAT is used with the neural network the training parameters applied are:

- Max. learning rate, $\beta_0 = 1.0$
- Max. training cycles per update, $C_{train} = 100$
- Size of data set, $N_{data} = 100$
- Max. network RMS error, $\epsilon_{max} = 0.001$

These values have been chosen so that the time taken for the resource models to be corrected will be sufficient for only a small number of imaging operations.

Scenario 1: Damaged Solar Array

In this scenario the power produced by the solar arrays drops suddenly by 22W two days into a 10-day simulation. This could have been caused by an unfortunate micrometeorite impact, radiation damage or component failure. Firstly NEAT is tested without any feedback of resource levels. The resulting progression of the battery state-of-charge is shown in Figure 8. At the 2-day event the model of solar power generation that NEAT is using departs from reality and so in this open-loop configuration the battery SOC rapidly falls far below the 85% minimum. For this reason the original NEAT algorithm regularly updates its estimates of resource availability using measurements. With this basic feedback in place the performance is improved (see Figure 9) but the battery SOC is still too low much of the time.

The same scenario was run again but with the neural network learning included. As can be seen from Figure 10 the management of the battery SOC is greatly improved with the minimum 85% battery SOC limit being effectively enforced after a period of learning (approximately 1.5 days) is largely determined by the progressive replacement of older training data with data points reflecting the new solar array performance. Here the neural network has enabled autonomous operations to continue despite a severe drop in the supply of power.

Scenario 2: Faulty Antenna

Here the effect of a greatly reduced downlink data rate is considered. The same 10-day scenario is used but in this case the downlink data rate is reduced by a factor of 10, due to a simulated fault with the antenna two days into the simulation. With NEAT operating open-loop, without any feedback of resource information, this causes repeated failure of image takes as shown in Figure 11. These are caused by there being insufficient memory to store the images. Without any feedback and an incorrect model of the downlink data rate, NEAT inevitably overestimates the available memory and schedules too many images.

When the basic feedback normally used is introduced these failed image takes are greatly reduced because NEAT becomes aware of the lower than expected level of data storage capacity shortly after each downlink. There are still some failures (see Figure 12) when an image capture is commanded immediately after a downlink and so before NEAT has received feedback that there is less available memory than expected. Furthermore, NEAT continues to make future plans based on what is now a hugely over-optimistic model of downlink capacity, with imaging operations often being removed from the schedule at the last moment when feedback data arrives.

When the neural network is included it quickly learns about the reduced downlink data rate, as shown in Figure 13. Over the course of a day the old training data is replaced with new data showing the reduced data rate. After this period of adaptation the corrected model of the downlink ca-




Figure 8: Battery SOC without resource level feedback.




Figure 9: Battery SOC with basic resource feedback.




Figure 10: Battery SOC with neural network learning.

capacity closely follows the true value. Consequently there are no failed image captures in this test case, although it would not have been surprising if there had been a few whilst the neural network was still adapting.




Figure 11: Failed image captures after a fault develops with the antenna when no resource feedback is used.




Figure 12: Failed image captures after a fault develops with the antenna and basic resource feedback is used.




Figure 13: Neural network learning about a step change in the downlink data rate.

Conclusions

It is essential that planning algorithms maintain up-to-date estimates of resource levels, based on measurements, to prevent modelling errors accumulating. However, this will not reduce future modelling errors unless a learning element is introduced. This work has shown that a neural network can be used make resource use models adapt when the behaviour of a spacecraft changes. Continuously updating training data with measurements of true resource levels ensures that corrections to resource models are always up-to-date.

This approach has been incorporated into the NEAT algorithm and applied to the case of SSTL’s UK-DMC imaging satellite. The results show that the neural network was able to learn about changes in resource behaviour. This made planning more robust to faults developing on the spacecraft and enabled effective autonomous planning and scheduling to continue even after a sudden change in system performance.

Acknowledgements

This research was funded by the UK Defence Technology Centre in Point & Image Data Fusion as part of the Autonomous Decision Making project. Thanks to Surrey Satellite Technology Ltd. with their help relating to UK-DMC.

References

- Brambilla, A.; Lavagna, M.; Da-Costa, A.; and Finzi, A. E. 2005. Distributed planning and scheduling for space system flotillas. In *8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*.
- Carrel, A. R., and Palmer, P. L. 2005. An evolutionary algorithm for near-optimal autonomous resource management. In *8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*.
- Chien, S.; Rabideau, G.; Knight, R.; Sherwood, R.; Engelhardt, B.; Mutz, D.; Estlin, T.; Smith, B.; Fisher, F.; Barrett, T.; Stebbins, G.; and Tran, D. 2000. ASPEN - Automated space mission operations using automated planning and scheduling. In *International Conference on Space Operations (Space Ops 2000)*.
- Chu, V.; Da-Silva-Curiel, R. A.; Sun, W.; and Sweeting, M. 2000. Disaster monitoring constellation. In *51st International Astronautical Congress*.
- Fukunaga, A.; Rabideau, G.; Chien, S.; and Yan, D. 1997. Towards an application framework for automated planning and scheduling. In *4th International Symposium on Artificial Intelligence, Robotics and Automation in Space*.
- Funase, R., and Nakasuka, S. 2005. Cooperation between onboard executive and planner on ground and its demonstration on remote-sensing nano-satellite “PRISM”. In *8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*.
- Globus, A.; Crawford, J.; Lohn, J.; and Pryor, A. 2002. Scheduling earth observing satellites with evolutionary algorithms. In *NASA Scheduling Workshop*.
- Goldberg, D. E., and Lingle, R. 1985. Alleles, loci, and the travelling salesman problem. In *International Conference on Genetic Algorithms and Their Applications*.
- Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley.
- Hashisa, Y., and Palmer, P. L. 2001. Epicyclic motion of satellites about an oblate planet. *Journal of Guidance, Control, and Dynamics* 24(3):586–596.
- Lemai, S., and Ingrand, F. 2004. Interleaving temporal planning and execution for an autonomous rover. In *4th*

International Workshop on Planning and Scheduling for Space (IWSS04).

Lin, S. 1965. Computer solutions of the traveling salesman problem. *Bell Systems Technical Journal* 44:2245–2269.

Mai, Y., and Palmer, P. L. 2001. Fast algorithm for prediction of satellite imaging and communication opportunities. *AIAA Journal of Guidance, Control, and Dynamics* 24(6):1118–1124.

Rumelhart, D. E., and et al. 1986. *Parallel Distributed Processing*. MIT Press. 318–362.

Wolfe, W. J., and Sorensen, S. E. 2000. Three scheduling algorithms applied to the earth observing systems domain. *Management Science* 46(1):148–168.

Wu, S. F.; Brewer, A.; and Palmer, P. 2001. ImPredict: a fast image prediction software and its application in the SSTL off-axis image scheduling system. In *15th AIAA/USU Conference on Small Satellites*.