

# Human Frontal Lobes and AI Planning Systems

An interdisciplinary view of integrated planning and reaction

Richard Levinson

Recom Technologies

NASA Ames Research Center, AI Research Branch

Mail Stop: 269-2, Moffett Field, CA 94035

levinson@ptolemy.arc.nasa.gov

## Abstract

This paper presents an AI model of integrated planning and reaction based on neuropsychological theories of similar human behavior. The model forges a link between frontal lobe models and AI planning systems that can improve our understanding of situated planning through a combined perspective. This paper first sets the frontal lobes in context, and then summarizes the frontal lobe theories upon which the model is based. Those theories are then synthesized into a single computational model for an artificial frontal lobe, and the components are discussed from both neuropsychological and AI perspectives. The paper concludes with a discussion of evaluation issues.

## Introduction

Achieving goals autonomously in an uncertain and changing world can require both predictive and reactive capabilities. Although a few AI systems have been built that combine both predictive and reactive components [11, 9, 14], the integration of those components is not well understood. In order to understand the integration better, we are studying the neuropsychological basis for similar human behavior. At NASA, we are developing autonomous scientific instruments that integrate real-time experiment control with deliberative experiment planning [10]. However, the system may be either too reactive or too deliberative. How can we assess if an instrument has an appropriate balance between prediction and reaction so that it can safely be left on its own? The approach described in this paper is to develop and evaluate a self-regulation module based on theories of human frontal lobe function.

Human autonomy requires a fine balance between data-driven reaction and goal-driven deliberation. Routine reactions often become inappropriate, and must be replaced by planned responses in novel and unexpected situations. This capacity for deliberate action is often damaged in victims of localized frontal lobe damage due to strokes, tumors, or direct impact. One theory of frontal lobe function involves: anticipation, goal selection, planning, initiation, self-monitoring, and use of feedback [19]. Many victims of frontal lobe damage cannot regain independence in their life because they cannot substitute deliberate plans for conditioned reactions [18]. These patients

often score well on IQ tests and perform memorized routine tasks normally, but they cannot generate and execute novel behavior [12].

Neuropsychology provides theories for integrated planning and reaction, but the components are not very formal or specific. In contrast, AI provides concrete instances of planning and reaction modules, but the integration of those components is poorly understood. Thus, neuropsychology and AI bring complementary contributions to a common goal: understanding the nature of integrated planning and reaction. The model in this paper provides a mapping between frontal lobe theories and AI planning systems that facilitates an interdisciplinary approach toward this goal.

## Frontal Lobes in Context

We begin by setting the context for the discussion. *Neuropsychology* combines neurology and psychology in order to analyze the behavioral effects of localized brain damage. A.R. Luria, a pioneer in the field, proposed that the brain consists of the three functional units shown in Figure 1. The information in this section comes mostly from Luria's landmark study *The Working Brain* [13], with support from [15].




Figure 1: The three functional brain units

The first unit, the *subcortex*, is responsible for regulating general wakefulness, autonomic functions, and reactive motor skills. The second unit, called the *sensory cortex*, consists of the posterior region of the cortex, and is responsible for receiving, processing, and storing information. The third unit, called the *motor cortex* or *frontal lobes*, consists of the anterior (frontal) region of the cortex. This unit is responsible for programming, verification, and regulation of activity.

It is important to note the existence of two control loops in figure 1. Routine situations are processed subconsciously by the *reactive* control loop (solid arrows),

and novel situations are processed deliberately by the predictive control loop (hollow arrows). The shaded region indicates that the focus of this paper is on modules for execution, self-regulation and planning. Those modules are based on theories of frontal lobe function that are summarized in the next section. Although the motor cortex is the primary focus of this paper, complex perception and action requires the integration of all three units. Therefore we briefly describe the relevant aspects of the other two units.

The **subcortex** receives input from both the sensory and the motor units, and it regulates general wakefulness, autonomic functions, and reactive motor skills. General wakefulness is controlled by the *reticular activating system*, which functions like a power strip by providing activation levels to the different local systems of the brain. Involuntary and autonomic behavior is regulated by the *hypothalamus* and the fluidity, coordination and balance of reactive motor skills are maintained by the *basal ganglia* and *cerebellum*. Also, the *hippocampus* provides a subconscious sensory filter called the *orienting reflex*, that produces signals only when novel stimuli are perceived.

The **sensory cortex** provides *data-driven* processing that brings information *into* the brain. Sensory processing begins with raw auditory, visual, and body sensory signals arriving at their respective primary sensation areas. Those elements are then integrated in association areas for each sense modality which are then connected together in a tertiary region that provides multisensory integration.

The **motor cortex** (or **frontal lobes**) provides *goal-driven* processing that sends signals *out* from the brain. Motor signals begin in the tertiary *prefrontal* region which predicts, monitors, and evaluates the effects of reactive behavior, and intervenes when appropriate. Output from the prefrontal region is sent to the secondary *premotor* region which controls sequences of motor commands by sending signals to the primary *motor* region, which controls the individual effectors. Many connections to other brain regions indicate extensive frontal lobe influence over mental as well as physical activity. Some of these connections influence subcortical motor skills and activation levels, and others send anticipatory signals that inform the sensory cortex what to expect.

### Theories of Frontal Lobe Function

Our model is primarily a synthesis of the theories of Lezak [12], Sohlberg and Geyer [18], and Norman and Shallice [17]. We now summarize these theories to provide the rationale behind our model. These theories describe a consistent core of functionality where they overlap, but they also include some different elements, and are described from different perspectives.

Lezak focuses on issues of neuropsychological evaluation, and uses the term *executive functions* to describe four essential elements of frontal lobe function.

First, *Goal Formulation* is the ability to generate and select descriptions of desirable future states. Without this function, people simply do not think of anything to do. Second, *Planning* involves the selection of steps, elements, and sequences needed to achieve a goal. This requires the ability to recognize and evaluate choices. Third, *Carrying Out Activities* involves the ability to start, stop, maintain, and switch between planned actions. Impairment of this function can affect the execution of well defined plans without disturbing impulsive behavior. Fourth, *Effective Performance* involves the ability to monitor and repair activities.

Sohlberg and Geyer focus on the cognitive rehabilitation of frontal lobe functions. They developed a model based on problems observed with frontal lobe patients. The first component, *Selection and Execution of Cognitive Plans*, involves the ability to describe goals and procedures, to determine appropriate action sequences, to initiate activity, to repair plans, and to maintain persistent effort until a task is completed. The second component, *Time Management*, involves the ability to generate realistic schedules, and to perform the scheduled activities within given time constraints. The third component, *Self-regulation*, involves using feedback to control behavior. This requires an ability to inhibit internal and external impulses that would trigger inappropriate reactions. Impaired self-regulation produces *environmental dependency* when strong external stimuli trigger inappropriate reactive behavior. For example, one patient began to bake cookies whenever she saw an oven. Another common problem, *perseveration*, involves difficulty with stopping or switching between responses.

Norman and Shallice based their model of frontal lobe damage on an AI production system, specifically on Soar [9]. Figure 2 shows their model, in which the frontal lobes serve as a *Supervisory Attentional System* that biases a reactive production system.




Figure 2: The Supervisory Attentional System

The *perceptual* system triggers overlearned (routine) control schemata. Routine resource conflicts between competing schemata are resolved via *contention scheduling*, and the selected schemata are then executed. The *Supervisory Attentional System* (SAS) biases the selections of the contention scheduler in non-routine or *novel* situations,

Shallice and Burgess propose that frontal lobes are required to: (1) inhibit undesirable old responses, and (2) generate and execute desirable new responses. They argue that damage to the SAS can mimic (model) two major effects of frontal lobe damage. *Disinhibition*




Figure 3: An Artificial Frontal Lobe

will occur if the SAS fails to inhibit responses triggered by strong impulses, or *apathy* will occur if the SAS fails to generate new responses in the absence of strong triggering impulses. The contention scheduler is modeled after Soar's production rules, and like Soar, this system involves distinct modes for reactive and deliberate behavior. In this model, high-level programs are defined in terms of Schank's Memory Organization Packages (MOPs) [16]. A sketch of the SAS algorithm includes the following steps: detect a problem, generate a plan, set up MOP triggering conditions called markers, interrupt activity when markers are triggered, assess difficulty, replan, and execute MOPs.

### An Artificial Frontal Lobe

In this section, we present a computer model that was designed to combine and operationalize the above theories. First, let's define some terms. A *program* is a sequence of commands that produce effector signals. A program can be *conditioned* so that it is invoked whenever a specified sensory condition is perceived. Such a program that is triggered and executed without any deliberation is called a *default program* or a *routine*. Programs include *choice points* where subroutines, resources, and orderings must be selected. Although the program representation can vary, it should be hierarchically structured. Norman and Shallice support this requirement for hierarchy with their use of MOPs. Also, Sohlberg and Lezak each identify the tendency to get lost in details as a common result of frontal lobe damage. They stress that the planning process must operate at an abstract level. Thus, a pure state-space representation is likely to be problematic.

Figure 3 shows the components of our artificial frontal lobe, with solid and hollow arrows indicating the reactive and predictive control loops, respectively. The *execution* module triggers and executes routine programs "subconsciously". The *planning* module projects the triggered routines and generates rules that map planned conditions into execution advice for *starting, stopping, switching, continuing* and *modifying* routines. The *self-regulation* module provides the interface between planning and execution.

The self-regulation module consists of six "critics" that monitor routine execution to detect and manage events that require deliberate action. The *routine inhibition critics* block the execution of routines that are *irrelevant, ineffective, or interfering*. The *plan execution critics* apply the plan rules and repair the plan when necessary. The *goal-generator* identifies desirable future states as new goals and then initiates plan generation. These critics function as demons that operate in parallel and communicate asynchronously.

The planning and execution modules have been implemented with a system called PROPEL [11], but the self-regulation module has not yet been built. PROPEL's action representation is a dialect of LISP that includes nondeterministic assignment statements and subroutine calls. These *choice points* identify alternative resources and subroutines. PROPEL's planner and executor are tightly integrated because they use the same algorithms and data types for interpreting a shared program representation. The programs can either be projected (simulated) by the planner or executed in real-time using default heuristics. The planner produces rules that advise the real-time executor about choice point selections. These rules are based on Drummond's Situated Control Rules (SCRs) [4], but due to PROPEL's action representation, the preconditions of our rules describe a program control stack instead of pure state-space predicates. PROPEL's planner uses SCRs to advise the executor in the same way Norman and Shallice's Supervisory Attentional System biases their contention scheduler.

This paper focuses on a situated planner that is connected to its environment via real-time sensors and effectors. We are developing two applications. The first is a NASA application that autonomously plans and executes mineral analysis and microbiology experiments. The second domain is a tool to help victims of frontal lobe damage to maintain autonomy over their lives by helping them plan and execute daily activities such as getting lunch, doing the laundry, or going shopping. The following three sections further describe the execution, planning and self-regulation modules from both neuropsychological and AI perspectives.

## Execution

The function of the execution module is to produce reactive (real-time) behavior. Routines are triggered by matching their preconditions against sensory input within bounded time. We are using PROPEL for this module, but other reactive AI systems could also be used. PRS [6], ERE[1], XFRM[14] and Robo-Soar[9] all monitor sensor conditions to trigger pre-defined responses in real-time.

The **Routine-Trigger** selects programs whose preconditions are satisfied by sensory stimuli. This is like the contention scheduler in Norman and Shallice's model, which uses production rules to resolve routine conflicts between competing responses. We use mutually exclusive PROPEL program preconditions to implement this module. After being triggered, routines are passed to the **routine-executor**.

The **Routine-Executor** generates effector command sequences by following the steps of a given program. This function is similar to that of the premotor region of the frontal lobes. PROPEL provides this ability by using default heuristics to instantiate choice points and execute default programs in bounded time. As described earlier, complex routine behavior in humans involves subcortical regions which operate as low-level servo and homeostatic controllers. Some AI applications may require similar functionality such as obstacle avoidance and battery recharging routines.

## Planning

The **plan-generator** calls the **program-projector** to predict the effects of a default response and to explore alternatives. It then generates a set of *plan rules* that advise the **routine-executor** when to take deliberate action based on planned times and situations.

When the **Plan-Generator** receives a new goal, it first estimates the amount of available planning time and then enlists the **program-projector** to search for a plan. The available planning time estimate is initially based on previous execution durations, and then updated as planning and reasoning proceeds. When planning time runs out, plan rules are generated that advise the **routine-executor**. These rules are based on PROPEL's situated control rules (SCRs), but they have been extended to distinguish five different types of advice. The rules map projected times and conditions into execution advice to *start*, *stop*, *switch*, *continue*, or *modify* routines. Generated plan rules are passed to the self-regulation module's plan execution critics.

The **Program-Projector** is an AI planner that searches through a space of program instances by choosing between alternative programs, resources, and subroutine orderings. The projection process must be efficient and interruptable in order to meet real-time constraints. This module is also called by the routine inhibition critics to detect inappropriate routines. We are using PROPEL's planner for this module. It uses the same program interpreter as the **routine-executor**,

except at choice points, where it creates disjunctive program continuations as branches in its search space. Other AI planning systems that use hierarchical action representations such as O-Plan [2], SIPE[20], and XFRM [14] also provide the projection capability that is required for this module. Anytime planning methods [3] allow a planner to be interrupted in order to meet real-time constraints.

## Self-Regulation

The self-regulation module consists of six "execution critics" that monitor routine execution to detect and manage events that require deliberate action. Three *routine inhibition critics* block the execution of inappropriate routines, and two *plan execution critics* follow and repair the plans. The sixth critic, the **goal-generator**, identifies desirable future states as new goals and then passes them to the planner. Each of these critics model behavior that is universally associated with frontal lobe function, and they can also be applied to AI systems.

The **Routine Inhibition Critics** are informed whenever a routine is triggered. Routines will be inhibited if they are *irrelevant*, *ineffective*, or *interfering*. A routine is irrelevant if it achieves no active goal; it is ineffective if it will not achieve its intended effect; and it is interfering if it conflicts with pre-existing plans. Critics determine each of these properties by first estimating the amount of available planning time (as described above for the **plan-generator**), and then calling the **program-projector** for that amount of time.

The **Relevance** critic detects and inhibits triggered routines when they do not achieve any relevant (current) goal. Failure of this function in humans is called *environmental dependency* when strong external stimuli trigger irrelevant routines. For example, a distracted driver may select a routine exit instead of remaining on the highway as planned. This module also applies to reactive AI systems because environmental dependency can result when goal references are compiled out so that program triggering conditions depend only on external state conditions. Such data-driven routines require simulation by the **program-projector** to determine the relevance of their effects.

The **Effectiveness** critic detects and inhibits routines that would not achieve their intended effects in *novel*, *difficult*, or *dangerous* situations. Various methods may be used to evaluate effectiveness. The subcortical *orienting reflex* helps humans to detect novel situations nearly instantaneously by producing signals only in the presence of novel stimuli. Equally fast tests for machines include detecting when no responses are triggered, or when the preconditions of competing responses are not mutually exclusive (e.g. when Soar reaches an impasse or when multiple PRS procedures are enabled). In general however, costly search using the **program-projector** is required to detect more subtle conditions that may cause a default program to

fail. This is the way that AI systems like PROPEL, ERE, and XFRM typically integrate planning and reaction. For this purpose, Reaction-First Search algorithms [5] bias the planner to project the default program first. When a *relevant* routine is inhibited because it is ineffective, the *plan-generator* is invoked to plan a more effective response.

The *Interference* critic detects and inhibits routines that would interfere with pre-existing plans. Failure to inhibit those routines causes *distractions*. The *program-projector* is called to determine if the triggered routine would interfere with pre-existing plans. AI systems such as ERE, XFRM, and Robo-soar do not currently use their planners this way. However, this critic will apply when the systems operate for extended time periods and manage a diverse set of simultaneous goals, plans and routines. For example, a low-priority reaction may interfere with an unrelated but high priority plan. When a *relevant* routine is inhibited because it is interfering, the *plan-generator* is invoked to plan a non-interfering response.

The *Plan Execution Critics* are responsible for carrying out the plans produced by the *plan-generator*. The *Advisor* applies the plan rules, and the *Replanner* detects and corrects plan errors.

The *Advisor* matches the plan rules (SCRs) against the external and internal state in order to advise the *routine-executor* when to *start*, *stop*, *switch*, *continue*, and *modify* routines. This defines five types of SCR advice. The first four come from frontal lobe theory, and the fifth one comes from AI. Advice to *start* a program serves as a deliberate version of the *routine-trigger*. Failure to deliberately *stop*, or to *switch* between routines corresponds to perseveration or inflexibility in humans [12]. Advice to *continue* a routine until plan conditions are satisfied corresponds to persistence. These four advice types apply to AI systems when a routine's default start and stop conditions have been adjusted for a novel problem. Advice to *modify* a routine tells the *routine-executor* to make a non-default choice point selection. This is the original type of SCR advice found in PROPEL and ERE. Plan transformation methods like those used by XFRM could provide a different type of *modify* advice.

The *Replanner* detects and corrects plan errors. It compares expectations with observations to determine when plan assumptions fail. The expectations are based on the preconditions of the plan rules. The most common causes of plan failure are variant (nondeterministic) outcomes, exogenous events, and incomplete knowledge. In humans, the orienting reflex helps to detect some forms of surprise. Machines must also be able to react to surprises and replan when necessary. We intend to use dependency analysis [8] to identify and monitor plan assumptions and focus the planner's search based on asynchronous sensor reports. Any available planning time is used to repair the plan. PROPEL facilitates replanning because it uses the same

*interpreter* for both planning and execution. This allows the planner to evaluate error recovery options by processing the execution program control stack. Other AI work related to replanning includes SIPE's execution monitors, PRIAR's dependency analysis [8], and the transformational planning methods found in XFRM.

The *Goal Generator* triggers goal-driven activity when the execution system is inactive or based on asynchronous sensory processing. It applies reasoning and abstraction to sensory perception in order to generate and select desirable future states as new goals, which are then passed to the *plan-generator*. Humans and machines who cannot generate goals are passive and akinetic when no programs are triggered at all. There has not been much AI work on goal generation.

## Evaluation Issues

Each component of our model corresponds to a functional dimension based on neuropsychological theory. Systems that combine planning and reaction must possess competence in these dimensions in order to approach the level of autonomy associated with human independence. The model suggests the following evaluation dimensions: Can the system detect and inhibit irrelevant, ineffective, or interfering routines? Is it perseverative (inflexible) or environmentally dependent? Is it too easily distracted? Can it generate effective goals and plans? Can it start, stop, switch, continue, and modify routines according to plan? Can it detect and correct plan errors?

Neuropsychologists have learned some valuable lessons while trying to answer these questions. Evaluating the ability to generate behavior in novel situations is difficult because test formats impose routine structure on the behavior [12]. The subjective character of self-regulation also makes it difficult to objectively evaluate [18]. This made it necessary to develop a variety of quantitative, qualitative, and mixed evaluation techniques [18]. Our model provides a mapping that facilitates the adaptation of existing frontal lobe tests for use on AI planning systems. Although a complete study of the tests and how they can be adapted will not fit within this paper, the different approaches of Lezak [12], Sohlberg and Geyer [18], and Shallice [17] provide examples of some alternative methods.

Lezak describes many tests that range from standardized maze and block configuration puzzles to execution tests that require patients to produce a deliberate sequence of routine actions. Perseverative patients are unable to deliberately switch between the routine actions. A more qualitative test asks patients to make anything they want with a set of Tinkertoys, allowing them to independently plan and execute "a potentially complex activity". Sohlberg and Geyer use informal observations of the patient's weekly activities to rate impairment severity along the dimensions of their model. This naturalistic approach is designed

to place minimal testing constraints on patient behavior. Shallice combined two common AI domains. His test is like the Towers of Hanoi, except the towers contain red, green and blue beads instead of different size disks. The patients are required to achieve different goal states like those found in the Blocksworld.

Since our system is not fully implemented, we have not yet evaluated its performance. Our intention is to design the model dimensions into our application specifications. Application-specific limits on distractability, environmental dependency, perseveration and persistence will be requirements on our system. Due to the tricky nature of testing these dimensions, we expect to use a combination of quantitative and qualitative tests. Although some neuropsychological tests are more directly applicable than others, they suggest a variety of evaluation issues and methods. To gain confidence that a machine can truly operate autonomously, we may eventually want to compare it with human performance in comparable situations.

### Conclusion

We have presented a model for integrated planning and reaction based on neuropsychological theories of human frontal lobe function. The frontal lobe theories led to several extensions in our AI model of situated planning. The neuro notions of environmental dependency and distraction prompted methods for inhibiting irrelevant and interfering routines. These are nice extensions to the familiar AI method of inhibiting ineffective routines. Also, the neuro notions of initiation, perseveration and persistence led to the distinction of 5 types of SCRs that encode advice to start, stop, switch, continue and modify routine programs. The model also facilitates further neuropsychological exploration of frontal lobe function. The benefits of computer simulated brain damage have recently been discussed from both AI and neuropsychological perspectives in [7].

The apparent relation between human frontal lobes and AI planning systems warrants further study. Our model supports such study by providing a mapping between neuropsychological and AI models. This mapping supports the interchange of ideas and evaluation procedures for interdisciplinary research aimed at a common goal: understanding the nature of integrated planning and reaction.

### Acknowledgements

Thanks to Peter Robinson, David Thompson, John Bresina, Keith Swanson, Wayne Iba and Lisc Getoor for useful comments.

### References

- [1] Bresina, J. and Drummond, M. Integrating Planning and Reaction: A Preliminary Report. 1990. *Proceedings of the 1990 AAAI Stanford Spring Symposium*, Stanford, CA.
- [2] Currie, K., & Tate, A. 1991. O-Plan: the Open Planning Architecture. *Artificial Intelligence*, Vol 52, No. 1, Autumn 1991, North-Holland.

- [3] Dean, T., and Boddy, M. An Analysis of Time-Dependent Planning. *In Proc. of AAAI-88*, pp. 49-54, St. Paul, MN, 1988.
- [4] Drummond, M., Situated Control Rules. *Proc. of the First International Conference on Principles of Knowledge Representation and Reasoning*, pp. 103-113, Toronto, Canada, 1989.
- [5] Drummond, M., Swanson, K., Bresina, J., and Levinson, R., Reaction-First Search: Incremental Planning with Guaranteed Performance Improvement. *The proceedings of IJCAI-93*.
- [6] Georgeff, M.P. and Lansky, A.L., "Reactive Reasoning and Planning," *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, Seattle, Washington (July 1987).
- [7] Hinton, G., Plaut, D., Shallice, T. Simulating Brain Damage. *Scientific American*. October 1993.
- [8] Kambhampati, S., and Hendler, J., A Validation-structure-based theory of plan modification and reuse. 1992. *Artificial Intelligence*, Vol. 55, pp. 193-258.
- [9] Laird, J., and Rosenbloom, P., Integrating Execution, Planning and Learning in Soar for external environments. 1990. *The Proc. of the AAAI-90*. Boston, MA.
- [10] Levinson, R., Robinson, P., and Thompson, D. Integrated Perception, Planning and Control for Autonomous Soil Analysis. *Proceedings of the 9th IEEE Conference on AI for Applications*, Orlando, FL. 1993.
- [11] Levinson, R., A General Programming Language for Unified Planning and Control. *Artificial Intelligence special issue on Planning and Scheduling*. To appear 1994.
- [12] Lezak, M., *Neuropsychological Assessment*. Oxford University Press. New York, 1983.
- [13] Luria, A. R., *The Working Brain*. Basic Books Inc. New York. 1973.
- [14] McDermott, D. 1990. Planning Reactive Behavior: A Progress Report, in *Proceedings: Workshop on Innovative Approaches to Planning, Scheduling and Control*, San Diego, CA. (1990).
- [15] Nolte, J., *The Human Brain*, 1993. Mosby-Year Book Inc. St. Louis, MO.
- [16] Schank, R. C., *Dynamic Memory*, Cambridge University Press, Cambridge. 1982.
- [17] Shallice, T., Burgess, P., Higher-Order Cognitive Impairments and Frontal Lobe Lesions in Man. In *Frontal Lobe Function and Dysfunction*, edited by H. Levin, H. Eisenberg, and A. Benton. pp. 125-138. Oxford University Press, New York, 1991.
- [18] Sohlberg, M., Mateer, C., *Introduction to Cognitive Rehabilitation*. The Guilford Press, New York, 1989.
- [19] Stuss, D., Benson, D., *The Frontal Lobes*. Raven Press Inc. New York. 1986.
- [20] Wilkins, David., *Practical Planning: Extending the Classical AI Paradigm*. 1988. San Mateo, CA. Morgan Kaufman Publishers.