From	To	Reason	How
3	1	Got bored.	Stare at object a moment Stare at feet Start looking around again
3	5	Overseer came nearby	Look at object Glance at Overseer Look at object Freeze
4	Any	Get your act together, little Patient	Stop and stare a moment Blink, blink Shake head while looking down Big sigh Back to work
5	1	Overseer went away again; The coast is clear	Watch Overseer leave Sigh Turn away from Overseer Squash down Sigh again Look over shoulder at Overseer Turn back away from Overseer
6	Any but 5	The coast is clear I just made it up	Look carefully around Shake head at folly Sigh
6	5	I'm paranoid, but I was right!	Same as 1-5

FIGURE C.3: Mini-transitions for Explore World, continued

For **Mope By Fence**:
Low-level signifiers:

1. Look Out At World
2. Sigh
3. Walk Up and Down Fence

Transitions:

From	To	Reason	How
1	2	Life is bad! Wish I was out there!	Stop looking a moment Lost in reverie
1	3	Bored with spot Get better position	Look in the direction I am planning to walk. Focus on something there Walk, keeping eye on spot
2	1	I'm sad, but I still want to look	Interruption
3	1	Got to point where I can see the thing I want to look at	Turn to face and look at the thing intently

FIGURE C.4: Mini-transitions that make up Mope By Fence

For **Read Sign**:
Low-level signifiers:

1. Read line
2. React to line

Transitions:

From	To	Reason	How
1	2	Saw something interesting	Shock reaction or re-read
2	1	Mulled it over	Pause Return to reading

FIGURE C.5: Mini-transitions that make up Read Sign

For **Exercise**:
Low-level signifiers:

1. Bob up and down

No Transitions.

FIGURE C.6: Mini-transitions that make up Exercise

For **Head-Banging**:
Low-level signifiers:

1. Hit Head on Ground
2. Wait to See if Light Went Out

Transitions:

From	**To**	**Reason**	**How**
1	2	Wants to get light on	
2	1	Light went out again	Act surprised Try to get light on by shaking head
2	1	Light went out again	Show frustration by freaking out

FIGURE C.7: Mini-transitions that make up Head-Banging

For **Unknown Behavior**:
Low-level signifiers:

1. Sigh
2. Look Around
3. Watch Overseer

Transitions:

From	To	Reason	How
1	2	Get your act together, little Patient	See sigh transition for Explore World
1	3	Sigh reminds you of your evil enemy	Turn slowly to Overseer Sigh again.
1	3	Just remembered you should be scared; or Overseer came nearby	Whirl around.
2	1	This place is bad	should work as interruption
2	3	Notice Overseer	Glance at Overseer Double-take If nearby, tremble and back up
3	1	What a pathetic piece of lamphood	Look away. Sigh.
3	2	Overseer went away	sigh and pause

FIGURE C.8: Mini-transitions that make up Unknown Behavior

For **Be Killed**:
Low-level signifiers:

1. Fear City
2. Die

Transitions:

From	To	Reason	How
1	2	Overseer hit button	Lightning bolt flash

FIGURE C.9: Mini-transitions that make up Be Killed

For **Turned Off**:
Low-level signifiers:

1. Be Turned Off

No Transitions.

FIGURE C.10: Mini-transitions that make up Turned Off

Composing High-Level Signifiers with Maxi-Transitions

Once each of the high-level signifiers was implemented, it was time to combine them with maxi-transitions to form the complete Behavior of the Patient. The design step for this is similar to that of composing the low-level signifiers. At this step, each possible transition between high-level signifiers is considered. For each possible transition, I listed the reasons for that behavioral change and corresponding ways to communicate that reason to the user. Figures C.11 through C.19 show the maxi-transition design for the Patient's high-level signifiers.

From **In Monitor:**

New Behavior	**Reason**	**How**
Explore world	Stops being so scared	Become curious. Move towards front. Look around carefully. Hop out. Still be a little scared for a while.
Head-banging	I am broken!	Here it is important to be clear as to what is going on. The Patient should look surprised, shake its head. Maybe the Overseer should look in disgust. When light goes on, Patient should be happy again.
Unknown behavior	Going on too long (?)	Pass in object of interest.

FIGURE C.11: Maxi-transitions from In Monitor

From **Explore World:**

New Behavior	**Reason**	**How**
In Monitor	Overseer comes right as Patient is coming out	Look at Overseer. Scurry back into monitor.
Read Sign	Notices schedule in its wandering.	Glance at schedule while walking by. Look interested. Walk over to it.
Mope by Fence	Gets near fence. Is bumming (after exercise). Oh, outside world! How cruel you are!! Wish I was back there.	Start slowing down beforehand. Life is bad. Look out at world. Sigh.
Head-Banging	If this happens, it's because the light goes out.	Look shocked. Look at camera so user can see your light is out. Shake your little head. Sideways, up and down. Swings get wider. Smack that head.
Turned Off	You've been moving around too much, getting too excited. Overseer doesn't like that	Maybe with your back to the Overseer, all of the sudden slump down. If you do see the Overseer, get scared. But keep moving so user sees the contrast. It's Overseer's job to make clear this is because of it.
Unknown Behavior	Explore world is going on a long time (?)	Pass in object of interest.

FIGURE C.12: Maxi-transitions from Explore World

From **Read Sign:**

New Behavior	**Reason**	**How**
Explore World	Overseer didn't bother it for some reason and it is done reading.	Stop reading. Stare at schedule. Sigh. Turn around. Look at world. Start exploring.
Exercise	Overseer comes over.	Sequence of looking at schedule and Overseer. Getting intimidated into it. Slow down as Overseer goes away.
Mope by Fence	Overseer is gone. Life is bad.	Slow down and stop. Sigh. Mope a little. Look at outside world. Sneak to the fence. Start moping.
Head-Banging	Light is out.	Make it short. Interruption.
Turned Off	Didn't pay attention that it was supposed to exercise	Like transition to exercise, but too sad to exercise. Sigh while looking at Overseer. Some pathetic attempts to exercise.

FIGURE C.13: Maxi-transitions from Read Sign

From **Exercise:**

New Behavior	**Reason**	**How**
Explore World Mope by Fence	Overseer is gone, it is bored and sad.	Must be sneaky. Slow down. Start looking around. Stop. Look at Overseer. Sneak off in other direction while keeping an eye on the Overseer.
Read Sign	Not done examining sign yet. Exercise is boring, sign is more interesting.	You're near the sign anyway (check). Turn around and start reading again. But slow down because you're not paying attention to exercise.
Head-Banging	Light goes out.	Quick interruption
Turned-Off	Not exercising enthusiastically enough.	Exercise slowly. Don't notice Overseer coming. When Overseer comes near, exercise frantically, back up a little, but it's too late.
Unknown Behavior	I don't know	Object of interest

FIGURE C.14: Maxi-transitions from Exercise

From **Mope by Fence:**

New Behavior	**Reason**	**How**
Explore World	Bored of looking out of fence. Life must go on.	One last big sigh. Turn around. Scan Industrial Graveyard. Start exploring, but sadly.
Exercise	Supposed to be exercising. Overseer comes near.	Glance at Overseer. Turn back to fence. Slow exercises.
Head-Banging	Light goes out.	Look surprised (but resigned). Do a frustration dance.
Turned Off	Supposed to be exercising, but didn't notice Overseer.	Turn around at last second and cringe.
Unknown Behavior	I don't know	Object of interest

FIGURE C.15: Maxi-transitions from Mope by Fence

From **Head-Banging:**

New Behavior	**Reason**	**How**
Explore World Read Sign Exercise	Head-banging as interruption	If short: light goes back on right away, go back to activity.
Mope by Fence	Mope by fence is more serious	
Be Killed	Overseer noticed and is angry.	Uh-oh! When Overseer is near, start cringing. Look around, trying to figure out when Overseer is near, but can't see anything. Back up, maybe bumping into stuff.
Turned Off	Overseer saw and doesn't like it.	Patient doesn't notice Overseer coming. Just turn off (sudden break).
Unknown Behavior	?	?

FIGURE C.16: Maxi-transitions from Head-Banging

From **Be Killed**:
No transitions once you're dead.

FIGURE C.17: Transitions from Be Killed

From **Turned Off:**

New Behavior	Reason	How
Explore World	Turn off over	Slowly rise up. Shake self. Blink, blink. Maybe sigh. Look around slowly to get orientation. This should be exaggerated the first time, after that it becomes a routine.
Exercise	Same	Here you should be exercising like a maniac while looking around for the Overseer. Taper off.
Mope by Fence	Just another reason to be depressed	Same as first transition, but even more depressed.

FIGURE C.18: Maxi-transitions from Turned Off

From **Unknown Behavior:**

New Behavior	Reason	How
In Monitor	Overseer came near	Freak out and back up
Explore World	Bored of standing there	Fixate on a point; start walking towards there
Read Sign	You're near the sign anyway, and you haven't read it yet.	Glance at sign. Look with more interest. Start going.
Exercise	Overseer came near and it is time.	Look at Overseer. Look surprised. Go nuts.
Mope by Fence	Life is sad.	Sigh. Sweep your gaze across the inside of the junkyard. Look at the outside world. Then switch over whole-heartedly.
Head-banging	Light goes out.	Just like everyone else.
Turned off	Should be exercising.	See transition from mope by fence to head-banging.

FIGURE C.19: Maxi-transitions from Unknown Behavior

C.3 Complete Patient Design

Once these maxi-transitions are implemented, the Patient is complete. The full patient design is shown in Figure C.20. However, due to time constraints the entire design was not implemented. The design of the Patient as implemented is shown in Figure 7.54.

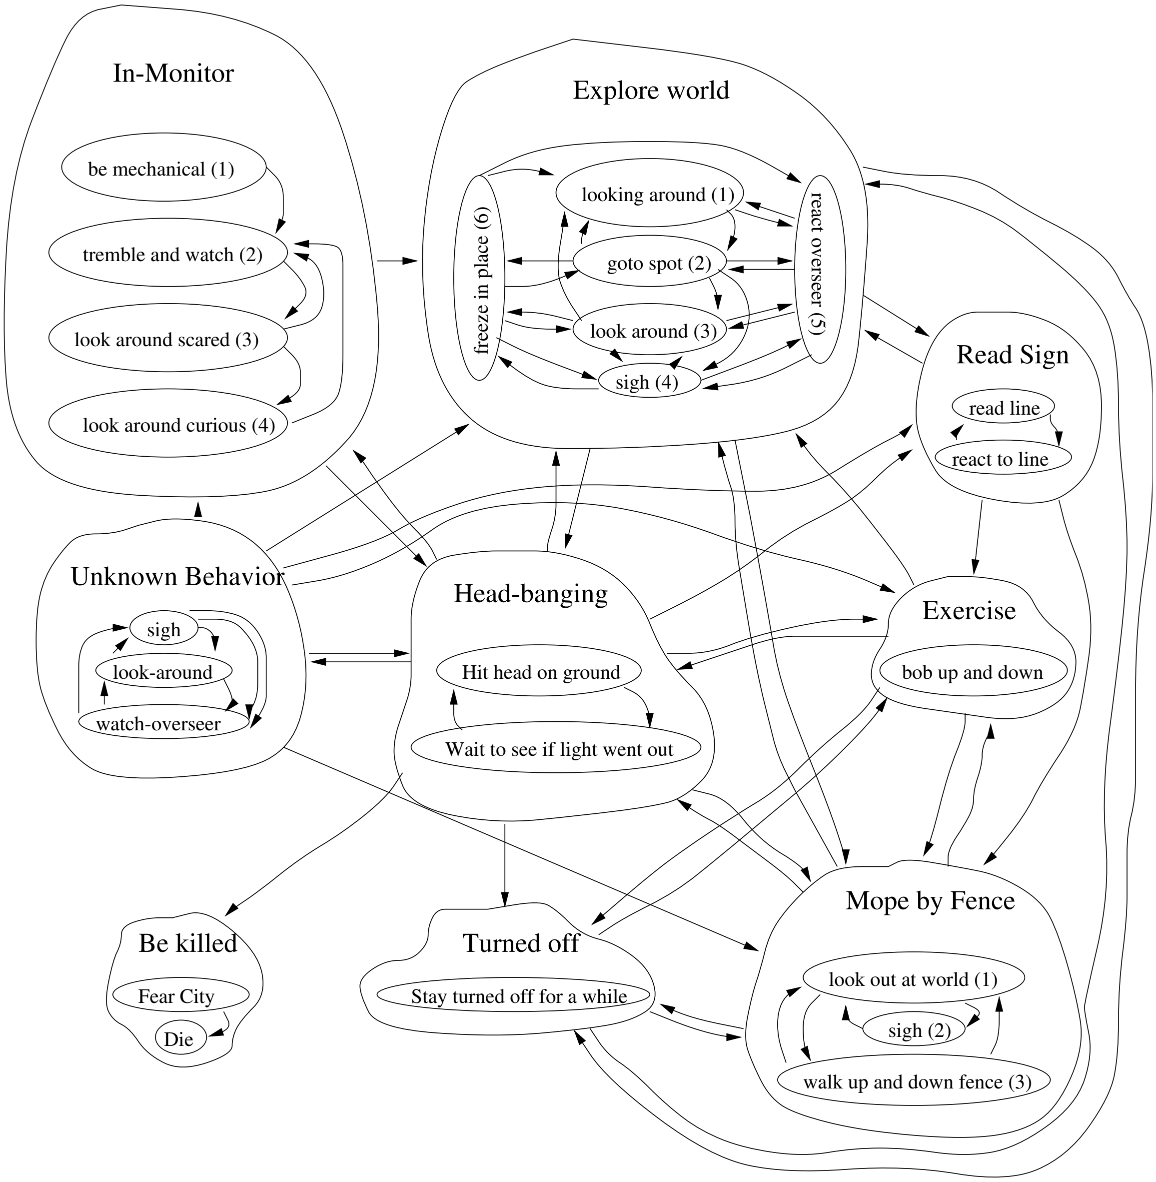

FIGURE C.20: The complete design of the Patient

Appendix D

Expostulations on Chapter 7 for the Technically Inclined

This Appendix consists of additions to Chapter 7 that are oriented for the reader whose technical interests are not exhausted by that populist rendering of the Expressivator. Pointers within the body of Chapter 7 will tell you when to read which part of this appendix.

D.1 Details on Transition Implementation

D.1.1 Transition Triggers

Transition triggers are complex sensors that look at conditions in the world to determine when it is time to switch from one behavior to another. Typically triggers test for things like the following:

- what behaviors are currently or have recently been run,
- what signs have recently been posted,
- events occuring in the virtual environment,
- communicative features,
- other transitions.

For example, when the Patient is hitting its head against the ground, it gets frustrated from time to time and switches from the "head-banging" low-level signifier to the "act frustrated" one. In order to determine when it is appropriate to switch, the transition trigger waits until the head-banging signifier has started running, and then counts the number of times a "smack head" sign has been posted, which corresponds to the number of times the user has seen the agent hit its head on the ground. After a sufficient number of smacks have occurred without the light going back on, the "act frustrated" transition trigger suggests to the rest of the system that it is time to change to the Patient's frustrated hopping-around behavior.

In general, when a trigger has found the right conditions for itself, it announces that fact to the rest of the system by finding the transition data structure associated with the signifier that called it, and modyifying it to reflect the trigger's opinion of what should be done. In particular, it notifies its parent signifier of which behavior should be terminated, which behavior should be started, and why.

D.1.2 Transition Demons

Transition demons keep an eye on the transition memory structures. They fire when an appropriate trigger has happened. Because it is generally more important to anticipate the new behavior properly than to finish up the old behavior in any particular way, transition demons generally check for transitions that are going to a particular behavior for a particular reason. Sometimes, they also check for the old behavior the agent was running.

The demon's job is to terminate the old behavior, go through a sequence of actions to create a transition, and then start the new behavior. The only exception is when the agent should merely interrupt a behavior, not terminate it; then the demon should make a transition, run the new behavior, and on termination make a transition back to the old behavior.

The transition demons' job is to kill the old behavior, do a transition sequence, and then start the new behavior. This sounds straightforward, but things are slightly more complicated. In particular, transitions must go from one behavior *that the user has seen* to another. For example, suppose the Patient has just decided to change from "look around scared" to "look around curiously." It has just killed off the "look around scared" behavior and is about to be curious when the Overseer approaches. Immediately, it is time for the Patient switch to "tremble and watch Overseer." Internally, this would mean a switch from "look around curious" to "tremble and watch Overseer" — but since the user does not know that the Patient is becoming curious, the correct transition is from "look around scared" to "tremble and watch Overseer." If this correctly chosen transition demon attempts to simple-mindedly kill the behavior from which it comes, "look around scared" (which no longer exists) will be killed and "look around curious" will continue on its merry way, running simultaneously with "tremble and watch Overseer." Oops.

To solve this kind of problem, transitions first delete, not just the old behavior they believe is running, but all other 'competing' behaviors. That is to say, mini-transitions kill any other low-level signifier that shares 'their' high-level signifier. Similarly, maxi-transitions kill any other high-level signifier. In order to make sure that now out-of-date transitions are deleted appropriately as well, mini-transitions are themselves declared as low-level signifiers, and maxi-transitions as high-level signifiers.

After transitions have killed preceding behaviors, they do some kind of transition sequence, and then start the requested new behavior. Mini-transitions add this to the high-level signifier that called them; maxi-transitions put it with the other high-level signifiers on the root of the agent's behavior tree.

If you have not yet had your fill of technical minutiae about transition implementation, I now refer you to section A.3 of the Appendix.

D.2 Technical Aspects to Expressivator Mindset Changes

The Expressivator was intended mainly as a way to add transitions to the basic Hap architecture, leaving the ordinary behavior structure alone. Nevertheless, it ends up fundamentally changing the meaning of both action-selection and behaviors in Hap.

D.2.1 Action-Expression in the Expressivator

Many agent architectures, especially those influenced by classical planning, require the agent designer to design behaviors based on their logical structure. For example, behaviors may be annotated with *preconditions* that state when they can be engaged in, and *postconditions* that note what changes they make to an environment. Action-selection then becomes a kind of problem-solving; you give the agent a goal to achieve in the world, and the agent chains behaviors until the last behavior's postcondition guarantees that the goal has been reached.

But there are many cases in which the 'point' of a behavior is not the changes the behavior may make in the environment, but the very behavior itself. 'Dancing,' for example, does not have any meaningful postconditions; the point of dancing is not to cause changes in the environment (unless it is a rain dance!), but for the pleasure of the activity itself[1]. The steps of the dance are not connected to one another by logical reasoning but by convention. There is, for example, no meaningful way for an agent to deduce that a foxtrot must consist of two long and two short steps; that's simply the way it's done. Many activities that are rooted in culture are similar. People usually do not stop for a rational analysis of when it is appropriate to say "hello," "thank you," or to ask someone how they are; they simply do it because it is conventional.

The action-selection mechanism in Hap is intended to reflect this concept of behaviors, not as means to achieve goals, but simply as sequences of actions to be engaged in for their own sake. Rather than having a designer specify the pre- and post-conditions for behaviors, both allowing and forcing the agent to reason about behavior before being able to act, the default in Hap is to have the designer specify behaviors as context-sensitive sequences of actions. The 'foxtrot' behavior will consist of the two long and two short steps simply because the designer wrote it that way; 'dancing' is done, not when the goal of dancing is achieved, but simply when the sequence of actions that make up dancing are complete.

[1] Of course, you can always call "pleasure" an effect on the world and measure the "pleasurableness" of various activities, and have the agent solve the problem of being happy by searching for and applying its maximally pleasurable activities. This neatly nails physical pleasure into the procrustrean bed of rationality by reducing temporally extended activity to a single step of goal satisfaction. There is little doubt that with sufficient ingenuity any activity can be mangled into goal-seeking rationality, but there should be at least a little scepticism about whether this is the best way of thinking about the problem.

In this framework, action-selection works largely by checking which behavior the agent is currently running and choosing the next step in the behavior. Behaviors are not simply scripted; they include annotations that let the agent know when the behavior is meaningful and when a running behavior no longer makes sense. However, by and large the reasoning behind the behavior's structure is implicit in the code for the behaviors as written by the designer.

Unlike Hap, the Expressivator demands that you know *why* your agent is doing what it does. The reasons for the agent's behavioral changes must be explicitly articulated in order to be expressed in transitions. Since transitions determine when and how it is appropriate to change from one signifier to another, they largely take over the role of action-selection for signifiers from the underlying architecture[2]. This means that, unlike Hap, explicit reasons for behavioral change form the basis for action-selection in the Expressivator.

At first, this change to Hap seemed unnatural: there did not seem to be any a priori reason why Hap action-selection should be inadequate for transitions. But the entire point of transitions is to show why you are switching from one behavior to another. If behaviors are simply sequenced, this means at some level you do not know why the behaviors are following one another; they simply do. That these reasons do not need to be articulated is an advantage in Hap because you do not always want to explain in fully logical, machine-understandable terms why the agent should do what it does. Nevertheless, it is a disadvantage if you want to express these reasons.

The Expressivator approach to action-selection is a compromise between the desire to include behavior whose logical structure cannot easily be elucidated, and the necessity to make reasons for behavioral choices explicit in order to express them. This is because the 'reasons' upon which the transitions are based need to be articulated *to the designer*, but not *to the machine*. Reasons for behavioral change are marked on transitions simply as tokens, such as "Patient-is-bored" or "Patient-saw-something-more-interesting." These reasons are not used by the agent to decide which activity makes sense, but by the designer as reminders of what the transition demon should express.

Still, the Expressivator does not include a full-fledged action-selection mechanism. For example, it could be that more than one transition triggers simultaneously, suggesting two conflicting behavioral changes. The Expressivator provides no mechanism to sort out which behavioral change should actually happen. I followed the style of Pengi, in making sure by hand that only one transition would ever fire in a particular circumstance. This strategy is not as painful as it sounds, because transitions are highly localized: (1) mini-transitions and maxi-transitions are handled separately; (2) mini-transitions can only ever conflict within their parent high-level signifier; (3) multiple transitions will only simultaneously fire when they are both transitions out of the same behavior. Nevertheless, it may be necessary in the future to add a full-fledged behavioral switching arbitration mechanism somewhat like that provided by Soar, which may check such things as the priorities of the various behaviors in question.

[2] Ordinary Hap action-selection still occurs for the subbehaviors which implement the signifiers

```
(parallel_production head_banging ()
  ;; initialize a pointer to myself
  (locals ("this_plan" "hh_plan_obj")
  ;; start my mini-transitions
  (with persistent
        (priority_modifier 200)
        (subgoal hit_head_to_wait_demon1 $$this_plan))
  (with persistent
        (priority_modifier 200)
        (subgoal hit_head_to_wait_demon2 $$this_plan))
  (with persistent
        (priority_modifier 200)
        (subgoal wait_to_hit_head_demon $$this_plan))
  (with persistent
        (priority_modifier 200)
        (subgoal freak_out_then_hit_head_demon
                 $$this_plan))
  ;; initialize my low-level signifiers
  (with (priority_modifier 100)
    (subgoal init_lls_headbanging))
  ;; start the first low-level signifier
  (with low_level_signifying
        (subgoal wait_for_light_on))
  ;; wait until the user notices me so I can
  ;; post myself
  (with effect_only
      (priority_modifier 300)
      (demon
        (("G (Goal CurrentLowLevelSignifier
                       == do_headbanging;);"))
        (post_high_level_signifier head_banging)))
  (wait))
```

FIGURE D.1: How Headbanging is invoked.

D.2.2 Behaviors in the Expressivator

High-level behaviors in Hap are simply some (context-sensitive) sequence of actions. In the Expressivator, on the other hand, a high-level signifier has a pre-given structure. Specifically, a high-level signifier consists simply of a set of low-level signifiers and the mini-transitions between them. When a high-level signifier is invoked, it simply starts a set of transition triggers and demons and the first low-level signifier (an example is in Figure D.1). After that, changes in low-level behaviors occur automatically as transitions trigger and then are implemented by transition demons.

Similarly, the full activity of the agent consists of the high-level signifiers and the maxi-transitions between them. When the agent starts up, it invokes all the maxi-transitions, and then starts the first high-level signifier. After that, the transitions take care of all subsequent changes to the agent's activity, triggering changes and modifying the agent's

The head-banging signifying behavior does the following things:

1. start its 4 mini-transitions
2. initialize its low-level signifiers
3. start the first low-level signifier
4. wait for the user to notice that it is happening, and then post itself to general memory

FIGURE D.2: Translation of previous figure for non-agent-builders and other interested parties.

behavioral structure as appropriate.

D.3 Behavior Transition Types

D.3.1 Explanatory Transition

The explanatory transition was the most useful, and I ended up using it for the majority of the transitions. They are easy to write — basically, you just make a short sequence of actions to explain what the agent is doing. Most of the time, they worked well. The only problem with explanatory transitions is that if you spend a lot of time in the explanatory sequence, the agent becomes less reactive. For example, the agent may be busy showing the user why it is about to read the schedule, and therefore not notice that the Overseer is about to attack. This problem can be ameliorated by varying the priority of various transitions, so that in this example the transition to reacting to the Overseer takes over even if the Patient is already in mid-transition. But in general, I found it was best to try to keep the transitions relatively short, if necessary by using meta-level controls to graft a transition-related activity onto the next behavior instead of doing it in the transition itself.

D.3.2 Subroutine Behavior Blend

A subroutine behavior blend involves combining two behaviors by adding a subroutine of the first behavior to the second behavior. For example, when the Patient goes from trembling at the Overseer to looking around scared, this is implemented by adding glances at the Overseer to look around scared. The subroutine behavior blend was easy to implement and did not require a lot of debugging. On the other hand, it was not so helpful from a narrative point of view; the behaviors probably would have made more sense with a clear, explained break between them.

The Mystery Transition

Relatively frequently, I would add a subbehavior to the new behavior, but it was not actually part of the old behavior, so it is not an ‘official’ subroutine behavior blend. For example, when the agent starts hitting its head in headbanging, the transition starts the headbanging behavior and then adds to it a subbehavior to first shake its head a few times to show the user its light has gone out and it is trying to get it back on again. This works nicely, though it could also be implemented as an explanatory transition. The main advantages over doing the additional subbehavior instead of an explanatory transition are (1) it can blend in with the other new behaviors' subbehaviors and (2) it makes sure the agent knows that this is "really" part of the second behavior, i.e. the current low-level signifier is set correctly as the new behavior instead of having the agent think it is in the nether region between the two behaviors.

D.3.3 Sudden Break

.

When used appropriately — i.e. not all the time, like in current architectures — this is both easy to do and very effective. The sudden break shows that the agent is having a visceral response to something going on around it. For example, when the Overseer comes near the patient, there is often a sudden break as the Patient whirls to face the Overseer and start trembling. Making this a sudden break makes it clear the Patient is not cogitating on the subject of the Overseer but rather having an immediate and intense reaction to it.

D.3.4 Interrupt

I use the interrupt-style transitions for behaviors that erupt during other behaviors. For example, the Patient may interrupt itself to sigh, and being turned off is also an interruption.

In general, I think the interrupt is dangerous. The turned off behavior, for example, can last a long time, and you probably don't want to return directly to the part of the behavior you were in last. For example, after being turned off by the Overseer and waking back up again, the Patient probably should not look intently at exactly the same spot on the trash in the world that it was looking at before.

This problem is compounded in Hap by the fact that behaviors don't really have any way of telling when they were interrupted (though signifiers could figure it out by seeing if they are still posted). This means after returning from an interrupt, a behavior may never realize anyone interrupted it; the behavior is completely oblivious to a fact that is essential to the user. In general, I think it would be better for behaviors like turned off that last a long time to kill the old signifier and start it all over again when they are done.

D.3.5 Reductive Behavior Blend

The reductive behavior blend reduces one of the behaviors to an attribute whose value can vary over time. The attribute is then applied to the other behavior. For example, when the Patient goes from looking around scared to looking around curiously, it first spends some time doing the scared version with fear set to a low value. Then, it goes to curious. This was easy to implement and blended the behaviors well: you could not tell when the change came between the scared behavior and the curious behavior. But for the same reason, this is a bad transition type from a narrative point of view: you do not know the agent is actually changing behaviors, or why the agent is becoming less scared. Again, a clear break with an explanation in between might have been more effective.

D.3.6 Off-screen Transition

I use this for almost all my transitions — the offscreen transition is built into the compiler. If the user is not looking at the agent, it immediately switches to the next behavior without a transition. This is useful in my system because transitions represent a kind of in-between state where the system is not totally sure which behavior it is in. It is therefore clearly best for the system to spend as little time in the transitions as possible.

This kind of transition might also be important in systems that have a function besides story or entertainment. In such a situation, it may be that transitions are for explanation, whereas the agent also has tasks to fulfill. In this case it's clearly best not to bother with explanation when the user is not paying attention.

For a few behaviors (for example, fear city to die) I left the transition in even when the behavior is not being watched. This was because the transition is so long that even if the user is not watching initially they may catch the end of the transition, and the transition is important enough to give the user the opportunity to see it. In some cases, I wait to change behaviors until I know the user is looking, so that s/he will not miss an important behavioral change.

D.3.7 Unknown Behavior

The unknown behavior is supposed to represent the default activity the agent does when it is not sure what to do. I wrote an Unknown Behavior for the patient, but I didn't end up using it in the system. If all your transitions are from and to a particular behavior, it doesn't make much sense to go to the unknown one for no reason. I also had a hard time coming up with good transitions for the Unknown Behavior since, by definition, you don't know why the agent is doing it. I therefore could not figure out how to get incorporate the unknown behavior in a logical way. It might be that in a different story — for example, where attention is not always focused on one agent — it may make more sense.

D.3.8 Principled Subroutine Behavior Blend

The idea of the principled subroutine behavior blend is to create a new behavior by combining already-running or new behaviors into a transition behavior. I use the principled subroutine behavior blend to go from being in the monitor to exploring the world. In this case, the Patient does a scared intermediate behavior that combines reacting to the overseer with stepping into the world while freezing in place at regular intervals.

This transition was difficult to write because it was basically like adding a whole new behavior. I could recycle some of the mini-transition demons but I also had to write some new ones specific to this maxi-transition. On the other hand, the behavior works well and is nicely reactive. In general, it is too much work, but it could be useful from time to time.

D.3.9 Symbolic Reduction

Under symbolic reduction, one behavior is reduced to a simple sign or symbol and incorporated into the other. I use this kind of transition when the Patient goes from reading the schedule to exercising. After launching the exercise behavior, I slowly reduce the energy as the Overseer goes away. This was very easy to write and works well. People definitely seem to understand what is going on.

D.3.10 Virtual Behavior Blend

In the virtual behavior blend, both behaviors run, but one of them has its muscle commands paralyzed. I use the virtual behavior blend when the Patient is turned off. This way, it would still have emotional reactions to the Overseer approaching, but would not actually move.

I found this kind of behavior blend exceptionally difficult to control. It had two major problems. Firstly, the agent would leap back into its old behavior the minute turned-off stopped paralyzing it, causing very strange behavioral discontinuities. Secondly, it was difficult to paralyze absolutely everything that needed paralyzing, with the result that the agent would still move around even though it was lying passed out on the ground. I fiddled with this transition extensively to get it right, but in the end, it did not seem to bring enough advantages to make it worth the effort.

D.4 Problems with Using Hap for Transitions

The number one problem with using Hap as a basis for the Expressivator is that you cannot pass around behavior names as Hap variables. Hap variables can only be integers, and for various reasons that have to do with the details of Hap's implementation in RAL it was not possible to encode goal names in a straightforward way as integers. The difficulty with this is that the transition system does some minimal reasoning about behaviors, and as soon as you start reasoning about them you need to be able to save them as variables. This would let you, for example, pass the behavior name to subbehaviors, save it in memory and call it later, and so on. Yes, it was always possible to find hacks around this problem, but this meant every instance of wanting to pass behavior names became an hour-long experiment in generating really horrific code.

This particular "feature" of Hap explains why I did write transitions that would go from any behavior to a particular behavior, but I never wrote transitions that went from a particular behavior to any other behavior. In order to do this, I would need to pass in to the generic transition the name of a behavior that it was going to have to start. But since I couldn't pass in the name of the behavior, this didn't happen. In general, I probably could have made a lot of the code much more general if I could have passed behavior names around.

Bibliography

[Agre and Chapman, 1987] Philip E. Agre and David Chapman. Pengi: An implementation of a theory of activity. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, July 1987.

[Agre and Chapman, 1990] Philip E. Agre and David Chapman. What are plans for? In Pattie Maes, editor, *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, pages 17–34. MIT Press, Cambridge, MA, 1990.

[Agre, 1988] Philip E. Agre. *The Dynamic Structure of Everyday Life*. PhD thesis, Massachusetts Institute of Technology Artificial Intelligence Laboratory, Cambridge, MA, 1988.

[Agre, 1990] Philip E. Agre. Portents of planning: A critical reading of the first paragraph of Miller, Galanter, and Pribram's 'Plans and the Structure of Behavior'. Paper presented at the Conference on Narrative in the Human Sciences, University of Iowa, July 1990., 1990.

[Agre, 1995] Philip E. Agre. The soul gained and lost: Artificial intelligence as a philosophical project. *Stanford Humanities Review*, 4(2), 1995. Special issue — Constructions of the Mind: Artificial Intelligence and the Humanities.

[Agre, 1997] Philip E. Agre. *Computation and Human Experience*. Cambridge University Press, Cambridge, UK, 1997.

[American Psychiatric Association, 1980] American Psychiatric Association. *Diagnostic and Statistical Manual of Mental Disorders (DSM-III)*. APA, Washington, D.C.:, 3rd edition, 1980.

[American Psychiatric Association, 1994] American Psychiatric Association. *Diagnostic and Statistical Manual of Mental Disorders (DSM-IV)*. APA, Washington, D.C.:, 4th edition, 1994.

[Analytix Inc., 1996] Analytix Inc. http://www.analytix.com/, 1996.

[Baker and Matlack, 1998] Stephen Baker and Carol Matlack. Chernobyl: If you can make it here... *Business Week*, pages 168–170, March 30 1998.

[Barthes, 1984] Roland Barthes. From work to text. In Brian Wallis, editor, *Art After Modernism: Rethinking Representation*, pages 169–174. New Museum of Contemporary Art, New York, 1984.

[Barton, 1995] Will Barton. Letting your self go: Hybrid intelligence, shared cognitive space and posthuman desire. Presented at Virtual Futures, Coventry, UK., 1995.

[Bates *et al.*, 1992] Joseph Bates, A. Bryan Loyall, and W. Scott Reilly. Integrating reactivity, goals, and emotion in a broad agent. Technical Report CMU-CS-92-142, Carnegie Mellon University, 1992. Also appeared in the Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society, Bloomington, Indiana, July 1992.

[Bates, 1994] Joseph Bates. The role of emotion in believable agents. Technical Report CMU-CS-94-136, Carnegie Mellon University, 1994. Also appears in Communications of the ACM, Special Issue on Agents, July 1994.

[Baur, 1991] Susan Baur. *The Dinosaur Man: Tales of Madness and Enchantment from the Back Ward*. Edward Burlingame Books, New York, 1991.

[Benyus, 1992] Janine M. Benyus. *Beastly Behaviors*. Addison-Wesley, Reading, MA, 1992.

[Blair and Meyer, 1997] David Blair and Tom Meyer. Tools for an interactive virtual cinema. In Robert Trappl and Paolo Petta, editors, *Creating Personalities for Synthetic Actors: Towards Autonomous Personality Agents*, number 1195 in Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 1997.

[Blanchot, 1981] Maurice Blanchot. *The Madness of the Day*. Station Hill, Barrytown, New York, 1981. Trans. Lydia Davis.

[Bledsoe, 1986] Woody Bledsoe. I had a dream: AAAI presidential address. *AI Magazine*, 7(1):57–61, 1986.

[Blumberg and Galyean, 1995] Bruce Blumberg and Tinsley A. Galyean. Multi-level direction of autonomous creatures for real-time virtual environments. In *Proceedings of SIGGraph*, 1995.

[Blumberg, 1994] Bruce Blumberg. Action-selection in Hamsterdam: Lessons from ethology. In *Proceedings of the 3rd International Conference on the Simulation of Adaptive Behavior*, Brighton, 1994.

[Blumberg, 1996] Bruce Blumberg. *Old Tricks, New Dogs: Ethology and Interactive Creatures*. PhD thesis, MIT Media Lab, Cambridge, MA, 1996.

[Brooks and Stein, 1993] Rodney A. Brooks and Lynn Andrea Stein. Building brains for bodies. Memo 1439, MIT AI Lab, August 1993.

[Brooks, 1986a] Rodney Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2:14–23, April 1986.

[Brooks, 1986b] Rodney A. Brooks. Achieving Artificial Intelligence through building robots. Memo 899, MIT AI Lab, Month 1986.

[Brooks, 1990] Rodney A. Brooks. Elephants don't play chess. In Pattie Maes, editor, *Designing Autonomous Agents*. MIT Press, Cambridge, MA, 1990.

[Brooks, 1991a] Rodney Brooks. Integrated systems based on behaviors. In *Proceedings of AAAI Spring Symposium on Integrated Intelligent Architectures*, Stanford University, March 1991. Available in *SIGART Bulletin*, Volume 2, Number 4, August 1991.

[Brooks, 1991b] Rodney A. Brooks. Intelligence without reason. Technical Report AI Memo 1293, MIT AI Lab, 1991.

[Brooks, 1994] Rodney A. Brooks. Coherent behavior from many adaptive processes. In Dave Smith, editor, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 22–29, Cambridge, MA, 1994. MIT Press.

[Brooks, 1995] Rodney A. Brooks. Intelligence without reason. In *The Artificial Life Route to Artificial Intelligence*, pages 25–81. Lawrence Erlbaum Associates, Hillsdale, NJ, 1995.

[Brooks, 1997] Rodney A. Brooks. From earwigs to humans. *Robotics and Autonomous Systems*, 20(2-4):291–304, June 1997.

[Bruner, 1986] Jerome Bruner. *Actual Minds, Possible Worlds*. Harvard University Press, Cambridge, MA, 1986.

[Bruner, 1990] Jerome Bruner. *Acts of Meaning*. Harvard University Press, Cambridge, MA, 1990.

[Bruner, 1991] Jerome Bruner. The narrative construction of reality. *Critical Inquiry*, 18(1):1–21, 1991.

[Carbonell, 1979] Jaime Carbonell. *Subjective Understanding: Computer Models of Belief Systems*. PhD thesis, Yale University Department of Computer Science, New Haven, CT, 1979.

[Chapman and Agre, 1986] David Chapman and Philip E. Agre. Abstract reasoning as emergent from concrete activity. In Michael P. Georgeff and Amy L. Lansky, editors, *Reasoning about Actions and Plans*, pages 411–424. Morgan Kauffman, Los Altos, CA, 1986.

[Chapman, 1990] David Chapman. *Vision, Instruction, and Action*. PhD thesis, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 1990.

[Cognitive Science, 1993] Cognitive science, January - March 1993. Volume 17. No. 1. Special Issue on Situated Action.

[Coleman *et al.*, 1984] James C. Coleman, James N. Butcher, and Robert C. Carson. *Abnormal Psychology and Modern Life*. Scott, Foresman and Co., Glenview, Illinois, 7 edition, 1984.

[Cooper, 1967] David Cooper. *Psychiatry and Anti-Psychiatry*. Ballantine Books, New York, NY, 1967.

[Crawford, 1993] T. Hugh Crawford. An interview with Bruno Latour. *Configurations*, 1(2):247–268, 1993.

[Crouch, 1990] Martha L. Crouch. Debating the responsibilities of plant scientists in the decade of the environment. *The Plant Cell*, pages 275–277, April 1990.

[Culler, 1982] Jonathan Culler. *On Deconstruction: Theory and Criticism after Structuralism*. Cornell University Press, Ithaca, NY, 1982.

[Curcuru,] Steve Curcuru. Personal Communication.

[Dautenhahn and Nehaniv, 1998] Kerstin Dautenhahn and Chrystopher Nehaniv. Artificial life and natural stories. In *International Symposium on Artificial Life and Robotics (AROB III)*, volume 2, pages 435–439, Beppu, Oita, Japan, 1998.

[Dawkins, 1989] Richard Dawkins. *The Selfish Gene*. Oxford University Press, Oxford, 1989.

[de Mul, 1997] Jos de Mul. Networked identities: Human identity in the digital era. In Michael B. Roetto, editor, *Proceedings of the Seventh International Symposium on Electronic Art*, pages 11–16, Rotterdam, 1997.

[DeLanda, 1991] Manuel DeLanda. *War in the Age of Intelligent Machines*. Zone Books, NY, 1991.

[Deleuze and Guattari, 1977] Gilles Deleuze and Félix Guattari. *Anti-Oedipus: Capitalism and Schizophrenia*. Viking Press, NY, 1977. Translated by Mark Seem.

[Deleuze and Guattari, 1987] Gilles Deleuze and Félix Guattari. November 28, 1947: How do you make yourself a body without organs. In *A Thousand Plateaus: Capitalism and Schizophrenia*, chapter 6, pages 149–166. University of Minnesota Press, Minneapolis, 1987. Translated by Brian Massumi.

[Dennett, 1987] Daniel Dennett. *The Intentional Stance*. MIT Press, Cambridge, MA, 1987.

[Derrida, 1976] Jacques Derrida. *Of Grammatology*. Johns Hopkins University Press, Baltimore, 1976. Translated by Gayatri Chakravorty Spivak.

[Don, 1990] Abbe Don. Narrative and the interface. In Brenda Laurel, editor, *The Art of Human-Computer Interface Design*, pages 383–391. Addison-Wesley, Reading, MA, 1990.

[Doray, 1988] Bernard Doray. *From Taylorism to Fordism: A Rational Madness*. Free Association, London, 1988. Trans. David Macey.

[Elliott *et al.*, 1998] Clark Elliott, Jacek Brzezinski, Sanjay Sheth, and Robert Salvatoriello. Story-morphing in the affective reasoning paradigm: Generating stories semi-automatically for use with 'emotionally intelligent' multimedia agents. In Katia P. Sycara and Michael Wooldridge, editors, *Proceedings of the Second International Conference on Autonomous Agents*, New York, May 1998. ACM Press.

[Foerst, 1996] Anne Foerst. Artificial Intelligence: Walking the boundary. *Zygon, Journal for Religion and Science*, 31:681–693, 1996.

[Foerst, 1998] Anne Foerst. Cog, a humanoid robot, and the question of the image of god. *Zygon, Journal for Religion and Science*, 33:91–111, 1998.

[Foner, 1993] Lenny Foner. What's an agent, anyway? http://foner.www.media.mit.edu/people/foner/Julia/Julia.html, 1993. Published in a revised version in *The Proceedings of the First International Conference on Autonomous Agents (AA '97)*.

[Forgy, 1991] C. L. Forgy. *Rule-extended Algorithmic Language Language Guide*. Production Systems Technologies, Inc., 1991.

[Foucault, 1973] Michel Foucault. *Madness and Civilization*. Vintage Books, New York, 1973. Trans. Richard Howard.

[Frank *et al.*, 1997] Adam Frank, Andrew Stern, and Ben Resner. Socially intelligent virtual petz. In Kerstin Dautenhahn, editor, *Proceedings of the 1997 AAAI Fall Symposium on Socially Intelligent Agents*, Menlo Park, CA, November 1997. AAAI Press. AAAI Technical Report FS-97-02.

[Fujita and Kageyama, 1997] Masahiro Fujita and Koji Kageyama. An open architecture for robot entertainment. In W. Lewis Johnson, editor, *Proceedings of the First International Conference on Autonomous Agents*, pages 435–442, New York, February 1997. ACM Press.

[Gadamer, 1986] Hans-Georg Gadamer. *Truth and Method*. Crossroad, NY, 1986.

[Galyean, 1995] Tinsley Galyean. *Narrative Guidance of Interactivity*. PhD thesis, MIT Media Lab, June 1995.

[Georgeff and Ingrand, 1989] M.P. Georgeff and F.F. Ingrand. Decision-making in an embedded reasoning system. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, pages 972–978, Detroit, MI, 1989.

[Goffman, 1961] Erving Goffman. *Asylums: Essays on the Social Situaiton of Mental Patients and Other Inmates*. Anchor Books, Garden City, NY, 1961.

[Goldstein, 1995] Kurt Goldstein. *The Organism: A Holistic Approach to Biology Derived from Pathological Data in Man*. Zone Books, NY, 1995.

[Gove, 1986] Philip Babcock Gove, editor. *Webster's Third New International Dictionary of the English Language Unabridged*. Merriam-Webster, Springfield, MA, 1986.

[Gross and Levitt, 1994] Paul R. Gross and Normal Levitt. *Higher Superstitions: The Academic Left and Its Quarrels with Science*. Johns Hopkins University Press, Baltimore, 1994.

[Guattari, 1984] Félix Guattari. *Molecular Revolution*. Penguin Books, New York, 1984. Trans. Rosemary Sheed.

[Haraway, 1990a] Donna Haraway. A cyborg manifesto: Science, technology, and socialist-feminism in the late twentieth century. In *Simians, Cyborgs, and Women: The Re-Invention of Nature*, pages 149–181. Free Association, London, 1990.

[Haraway, 1990b] Donna Haraway. Situated knowledges: The science question in feminism and the privilege of partial perspective. In *Simians, Cyborgs, and Women: The Re-Invention of Nature*, pages 183–201. Free Association, London, 1990.

[Harding, 1994] Sandra Harding. Is science multicultural? challenges, resources, opportunities, uncertainties. *Configurations*, 2(2):301–330, 1994.

[Hayes *et al.*, 1994] Patrick J. Hayes, Kenneth M. Ford, and Neil Agnew. On babies and bathwater: A cautionary tale. *AI Magazine*, 15(4):15–26, 1994.

[Hayles,] N. Katherine Hayles. How we became posthuman: Virtual bodies in cybernetics, literature, and informatics. Unpublished Manuscript.

[Hayles, 1993] N. Katherine Hayles. The materiality of informatics. *Configurations*, 1(1):147–170, 1993.

[Hayles, 1997] N. Katherine Hayles. Narrative and the question of embodiment in scientific inquiry. Society for Literature and Science, 1997.

[Horswill, 1993] Ian Horswill. Polly: A vision-based artificial agent. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 824–829, Menlo Park, July 1993. AAAI Press.

[Hsu *et al.*, 1990] Feng-hsiung Hsu, Thomas S. Anantharaman, Murray S. Campbell, and Andreas Nowatzyk. Deep thought. In T. Anthony Marsland and Jonathan Schaeffer, editors, *Computers, Chess, and Cognition*. Springer Verlag, New York, 1990.

[James, 1998] Ellen James. *Her Protector*. Harlequin Superromance. Harlequin Books, Toronto, 1998.

[Janet, 1889] Pierre Janet. *L'Automatisme Psychologique: Essai de Psychologie Expérimentale sur les Formes Inférieures de l'Activité Humaine*. Ancienne Librairie Germer Baillière et Cie, Paris, 1889. Ed. Félix Alcan.

[Joachims *et al.*, 1997] Thorsten Joachims, Dayne Freitag, and Tom Mitchell. WebWatcher: A tour guide for the World Wide Web. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)*, August 1997.

[Johnson, 1997] W. Lewis Johnson, editor. *Proceedings of the First International Conference on Autonomous Agents*, NY, February 1997. ACM Press.

[Josephson and Josephson, 1962] Eric Josephson and Mary Josephson. Introduction. In Eric Josephson and Mary Josephson, editors, *Man Alone: Alienation in Modern Society*, pages 9–53. Dell Publishing, NY, 1962.

[Keller, 1985] Evelyn Fox Keller. *Reflections on Gender and Science*. Yale University Press, New Haven, 1985.

[Kennedy, 1989] Noah Kennedy. *The Industrialization of Intelligence: Mind and Machine in the Modern Age*. Unwin Hyman, Boston, 1989.

[Kirk and Kutchins, 1992] Stuart A. Kirk and Herb Kutchins. *The Selling of DSM: The Rhetoric of Science in Psychiatry*. A. de Gruyter, New York, 1992.

[Laing and Esterson, 1970] R. D. Laing and A. Esterson. *Sanity, Madness, and the Family*. Penguin Books, Ltd., Middlesex, England, 1970.

[Laing, 1960] R. D. Laing. *The Divided Self: An Existential Study in Sanity and Madness*. Penguin Books, Middlesex, England, 1960.

[Laird, 1991] John Laird, editor. *Proceedings of AAAI Spring Symposium on Integrated Intelligent Architectures*, March 1991. Available in *SIGART Bulletin*, Volume 2, Number 4, August 1991.

[Lanier, 1996] Jaron Lanier. My problem with agents. *Wired*, 4(11), November 1996.

[Lasseter, 1987] John Lasseter. Principles of traditional animation applied to 3D computer animation. In *Proceedings of SIGGRAPH '87*, pages 35–44. Association for Computing Machinery, 1987.

[Laurel, 1986] Brenda Laurel. Interface as mimesis. In *User-Centered System Design*, pages 67–85. Lawrence Erlbaum Associates, Hillsdale, NJ, 1986.

[Laurel, 1991] Brenda Laurel. *Computers As Theatre*. Addison-Wesley, Reading, MA, 1991.

[Lecocq, 1991] Sylvain Lecocq. Untitled (letter to his doctor). In John G. H. Oakes, editor, *In the Realms of the Unreal: "Insane" Writings*, pages 158–161. Four Walls Eight Windows, New York, 1991. Trans. Roger Cardinal.

[Lester and Stone, 1997] James C. Lester and Brian A. Stone. Increasing believability in animated pedagogical agents. In W. Lewis Johnson, editor, *Proceedings of the First International Conference on Autonomous Agents*, pages 16–21, NY, February 1997. ACM Press.

[Lester *et al.*, 1997] James C. Lester, Sharolyn A. Converse, Susan E. Kahler, S. Todd Barlow, Brian A. Stone, and Ravinder S. Bhogal. The persona effect: Affective impact of animated pedagogical agents. In *Proceedings of CHI '97*, pages 359–366, Atlanta, March 1997.

[Levins and Lewontin, 1985] Richard Levins and Richard C. Lewontin. *The Dialectical Biologist*. Harvard University Press, Cambridge, MA, 1985.

[Lewontin *et al.*, 1984] Richard C. Lewontin, Steven Rose, and Leon J. Kamin. *Not In Our Genes: Biology, Ideology, and Human Nature*. Pantheon Books, New York, 1984.

[Lewontin, 1991] R. C. Lewontin. *Biology as Ideology: The Doctrine of DNA*. Harper Perennial, New York, 1991.

[Lewontin, 1995] Richard C. Lewontin. À la recherche du temps perdu. *Configurations*, 3(2):257–265, 1995.

[Loyall,] A. Bryan Loyall. Personal communication.

[Loyall and Bates, 1991] A. Bryan Loyall and Joseph Bates. Hap: A reactive, adaptive architecture for agents. Technical Report CMU-CS-91-147, Carnegie Mellon University, 1991.

[Loyall and Bates, 1993] A. Bryan Loyall and Joseph Bates. Real-time control of animated broad agents. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, Boulder, Colorado, June 1993.

[Loyall, 1997a] A. Bryan Loyall. *Believable Agents: Building Interactive Personalities*. PhD thesis, Carnegie Mellon University, Pittsburgh, May 1997. CMU-CS-97-123.

[Loyall, 1997b] A. Bryan Loyall. Personality-rich believable agents that use language. In W. Lewis Johnson, editor, *Proceedings of the First International Conference on Autonomous Agents*, NY, February 1997. ACM Press.

[Lukács, 1971] Georg Lukács. Reification and the consciousness of the proletariat. In *History and Class Consciousness: Studies in Marxist Dialectics*. MIT Press, Cambridge, MA, 1971. Trans. Rodney Livingstone.

[Maes, 1989a] Pattie Maes. The dynamics of action-selection. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, volume 2, pages 991–997, Detroit, MI, August 1989. Morgan Kaufmann.

[Maes, 1989b] Pattie Maes. How to do the right thing. AI Memo 1180, MIT AI Laboratory, December 1989.

[Maes, 1990a] Pattie Maes, editor. *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*. MIT Press, Cambridge, MA, 1990.

[Maes, 1990b] Pattie Maes. Guest editorial: Designing autonomous agents. In Pattie Maes, editor, *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, pages 1–2. MIT Press, Cambridge, MA, 1990.

[Maes, 1990c] Pattie Maes. Situated agents can have goals. In Pattie Maes, editor, *Designing Autonomous Agents*, pages 49–70. MIT Press, Cambridge, MA, 1990.

[Maes, 1991] Pattie Maes. Behavior-based Artificial Intelligence. In Jean-Arcady Meyer and Stewart W. Wilson, editors, *From Animals to Animats 2*, pages 2–10, Cambridge, MA, 1991. MIT Press.

[Maes, 1993 1994] Pattie Maes. Modeling adaptive autonomous agents. *Artificial Life*, 1(1-2):135–162, Fall-Winter 1993-1994.

[Mahoney, 1980] Michael S. Mahoney. Reading a machine. In N. Metropolis, J. Howlett, and G.-C. Rota, editors, *A History of Computing in the Twentieth Century: A Collection of Essays*, pages 3–9. Academic Press, NY, 1980.

[Mahoney, 1997] Michael S. Mahoney. Software and the assembly line. Carnegie Mellon Software Engineering Institute Invited Talk, February 1997.

[Marx, 1967] Karl Marx. *Capital: A Critique of Political Economy*, volume I: "The Process of Capitalist Production". International Publishers, New York, 1967. Ed. Frederick Engels. Trans. Samuel Moore and Edward Areling.

[Massumi, 1992] Brian Massumi. *A User's Guide to Capitalism and Schizophrenia: Deviations from Deleuze and Guattari*. MIT Press, Cambridge, MA, 1992.

[Mateas, 1997] Michael Mateas. Computational subjectivity in virtual world avatars. In Kerstin Dautenhahn, editor, *Proceedings of AAAI-97 Workshop on Socially Intelligent Agents*, pages 87–92, 1997. Available from AAAI as Technical Report FS-97-02.

[McDermott, 1981] Drew McDermott. Artificial Intelligence meets natural stupidity. In John Haugeland, editor, *Mind Design: Philosophy, Psychology, Artificial Intelligence*, chapter 5, pages 143–160. MIT Press, Cambridge, MA, 1981.

[Mérő, 1990] László Mérő. *Ways of Thinking: The Limits of Rational Thought and Artificial Intelligence*. World Scientific, New Jersey, 1990. Edited by Viktor Mészáros. Translated by Anna C. Gősi-Greguss.

[Miedaner, 1981] Terrel Miedaner. The soul of the Mark III beast. In Douglas R. Hofstadter and Daniel C. Dennett, editors, *The Mind's I: Reflections on Self and Soul*, pages 109–115. Basic Books, Inc., New York, 1981.

[Minsky, 1988] Marvin Minsky. *The Society of Mind*. Simon and Schuster, New York, 1988.

[Minuchin *et al.*, 1978] Salvador Minuchin, Bernice L. Rosman, and Lester Baker. *Psychosomatic Families: Anorexia Nervosa in Context*. Harvard University Press, Cambridge, MA, 1978.

[Mumford, 1934] Lewis Mumford. *Technics and Civilization*. Harcourt, Brace, and Company, New York, 1934.

[Neal Reilly, 1996] Scott Neal Reilly. *Believable Social and Emotional Agents*. PhD thesis, Carnegie Mellon University, 1996. CMU-CS-96-138.

[Newell and Simon, 1972] Allen Newell and Herbert A. Simon. *Human Problem Solving*. Prentice-Hall, Englewood Cliffs, NJ, 1972.

[Newell, 1981] Allen Newell. The knowledge level. Technical Report CMU-CS-81-131, Carnegie Mellon University Department of Computer Science, 1981.

[Newell, 1990] Allen Newell. *Unified Theories of Cognition*. Harvard University Press, Cambridge, Massachusetts, 1990.

[Nilsson, 1984] Nils J. Nilsson. Shakey the robot. Technical Report 323, AI Center, SRI International, Menlo Park, CA, 1984.

[Noble, 1998] David F. Noble. Digital diploma mills: The automation of higher education. *First Monday*, 3(1), January 1998. http://firstmonday.dk/issues/issue3_1/noble/index.html.

[Norman, 1993] Donald A. Norman. Cognition in the head and in the world: An introduction to the special issue on situated action. *Cognitive Science*, 17:1–6, 1993.

[O'Neill-Brown, 1997] Patricia O'Neill-Brown. Setting the stage for the culturally adaptive agent. In *Proceedings of the AAAI Fall Symposium on Socially Intelligent Agents*, pages 93–97, Menlo Park, CA, 1997. AAAI Press.

[Payton *et al.*, 1992] David W. Payton, David Keirsey, Dan M. Kimble, Jimmy Krozel, and J. Kenneth Rosenblatt. Do whatever works: A robust approach to fault-tolerant autonomous control. *Journal of Applied Intelligence*, 2:225–250, 1992.

[Pearson *et al.*, 1993] Douglas J. Pearson, Scott B. Huffman, Mark B. Willis, John E. Laird, and Randolph M. Jones. Intelligent multi-level control in a highly reactive domain. In *Proceedings of the International Conference on Intelligent Autonomous Systems*, Pittsburgh, PA, 1993.

[Penny, 1995] Simon Penny. Living machines. *Scientific American*, September 1995. 150th Anniversary Issue.

[Penny, 1997a] Simon Penny. Embodied cultural agents at the intersection of robotics, cognitive science, and interactive art. In Kerstin Dautenhahn, editor, *Socially Intelligent Agents: Papers from the 1997 Fall Symposium*, pages 103–105, Menlo Park, 1997. AAAI Press. Technical Report FS-97-02.

[Penny, 1997b] Simon Penny. The virtualisation of artistic practice: Body knowledge and the engineering world view. *CAA Art Journal*, Fall 1997. Ed. Johanna Drucker. Special Issue on Electronic Art.

[Perlin, 1995] Ken Perlin. Real time responsive animation with personality. *IEEE Transactions on Visualization and Computer Graphics*, 1(1):5–15, March 1995.

[Petta and Trappl, 1997] Paolo Petta and Robert Trappl. Personalities for synthetic actors: Current issues and some perspectives. In Paolo Petta and Robert Trappl, editors, *Creating Personalities for Synthetic Actors: Towards Autonomous Personality Agents*, number 1195 in

Lecture Notes in Artificial Intelligence, pages 209–218. Springer Verlag, Berlin, 1997.

[Pixar, 1986] Pixar. Luxo, jr. (film), 1986.

[Porter, 1997] Tom Porter. Depicting perception, thought, and action in Toy Story. In *First International Conference on Autonomous Agents*, February 1997. Invited Talk.

[Rheingold, 1995] Howard Rheingold. Technology criticism, ethics, and you. *San Francisco Examiner*, November 1 1995. Also available in ZK Proceedings: Net Criticism, Volume 1, p. 110, http://www.desk.nl/ nettime/zkp/.

[Rhodes, 1996] Bradley James Rhodes. PHISH-nets: Planning heuristically in situated hybrid networks. Master's thesis, MIT Media Lab, 1996.

[Ritzer, 1993] George Ritzer. *The McDonaldization of Society: An Investigation into the Changing Character of Contemporary Social Life*. Pine Forge Press, Newbury Park, CA, 1993.

[Robear Jr., 1991] James Walter Robear Jr. Reality check. In John G. H. Oakes, editor, *In the Realms of the Unreal: "Insane" Writings*, pages 18–19. Four Walls Eight Windows, New York, 1991.

[Ronell, 1989] Avital Ronell. *The Telephone Book: Technology — Schizophrenia — Electric Speech*. University of Nebraska Press, Lincoln, 1989.

[Rouse, 1993] Joseph Rouse. What are cultural studies of science? *Configurations*, 1(1):57–94, 1993.

[Sacks, 1995] Oliver Sacks. Forward to Kurt Goldstein's *The Organism: A Holistic Approach to Biology Derived from Pathological Data in Man*. Zone Books, NY, 1995.

[Samuel, 1995] Arthur L. Samuel. Some studies in machine learning using the game of checkers. In Edward A. Feigenbaum, editor, *Computer and Thought*. AAAI Press, Menlo Park, 1995.

[Schempf, 1995] Hagen Schempf. BOA: Asbestos pipe-insulation removal robot sydtem. Technical Report CMU-RI-TR-95-15, Carnegie Mellon University Robotics Institute, Pittsburgh, PA, 1995.

[Sengers,] Nishka Sengers. Personal Communication.

[Sengers, 1998] Phoebe Sengers. Do the thing right: An architecture for action-expression. In *Proceedings of the Second International Conference on Autonomous Agents*, May 1998. To appear.

[Shakes *et al.*, 1997] Jonathan Shakes, Marc Langheinrich, and Oren Etzioni. Dynamic reference sifting: A case study in the homepage domain. In *Proceedings of the Sixth International World Wide Web Conference*, pages 189–200, 1997.

[Shklovsky, 1990] Victor Shklovsky. *Theory of Prose*. Dalkey Archive Press, Elmwood, Ill, 1990. Trans. Benjamin Sher.

[Simmons *et al.*, 1997] Reid Simmons, Richard Goodwin, Karen Zita Haigh, Sven Koenig, and Joseph O'Sullivan. A modular architecture for office delivery robots. In W. Lewis Johnson, editor, *Proceedings of the First International Conference on Autonomous Agents*, pages 245–252, NY, February 1997. ACM Press.

[Smithers, 1992] Tim Smithers. Taking eliminative materialism seriously: A methodology for autonomous systems research. In Francisco J. Varela and Paul Bourgine, editors, *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 31–47, Cambridge, MA, 1992. MIT Press.

[Snow, 1969] C. P. Snow. *The Two Cultures and the Scientific Revolution*. Cambridge University Press, London, 1969.

[Social Text, 1996] Social Text, April 1996. Volume 14, Nos. 1-2. Ed. Andrew Ross. Special Issue on the Science Wars.

[Sokal, 1996] Alan Sokal. A physicist experiments with cultural studies. *Lingua Franca*, pages 62–64, May-June 1996.

[Sontag, 1979] Susan Sontag. *Ilness as Metaphor*. Vintage Books, New York, NY, 1979.

[Steels and Brooks, 1995] Luc Steels and Rodney Brooks, editors. *The Artificial Life Route to Artificial Intelligence*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1995.

[Steels, 1994] Luc Steels. The Artificial Life roots of Artificial Intelligence. *Artificial Life*, 1(1-2):75–110, 1994.

[Stone, 1996] Brian A. Stone. Dynamically sequencing an animated pedagogical agent. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, pages 424–431, Portland, OR, August 1996.

[Strasser, 1982] Susan Strasser. *Never Done: A History of American Housework*. Pantheon Books, New York, 1982.

[Suchman, 1987] Lucy A. Suchman. *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge University Press, Cambridge, 1987.

[Swados, 1962] Harvey Swados. The myth of the happy worker. In Eric and Mary Josephson, editors, *Man Alone: Alienation in Modern Society*, pages 106–113. Dell Publishing, NY, 1962.

[Sycara and Wooldridge, 1998] Katia P. Sycara and Michael Wooldridge, editors. *Proceedings of the Second International Conference on Autonomous Agents*, NY, May 1998. ACM Press.

[Thomas and Johnston, 1981] Frank Thomas and Ollie Johnston. *Disney Animation: The Illusion of Life*. Abbeville Press, New York, 1981.

[Tyrell, 1993] Toby Tyrell. *Computational Mechanisms for Action Selection*. PhD thesis, University of Edinburgh, 1993.

[Van den Haag, 1962] Ernest Van den Haag. Of happiness and of despair we have no measure. In Eric and Mary Josephson, editors, *Man Alone: Alienation in Modern Society*, pages 180–199. Dell Publishing, NY, 1962.

[Varela *et al.*, 1991] Francisco J. Varela, Evan Thompson, and Eleanor Rosch. *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press, Cambridge, MA, 1991.

[Various, 1985] Various. *Die Bibel*. Deutsche Bibelgesellschaft, Stuttgart, 1985. Translated by Martin Luther.

[Vere and Bickmore, 1990] Steven Vere and Timothy Bickmore. A basic agent. *Computational Intelligence*, 6:41–60, 1990.

[Viscott, 1972] David S. Viscott. *The Making of a Psychiatrist*. Arbor House, New York, NY, 1972.

[Walter, 1963] W. Grey Walter. *The Living Brain*. W.W. Norton, 1963.

[Washington, 1991] Karoselle Washington. The killing floors. In John G. H. Oakes, editor, *In the Realms of the Unreal: "Insane" Writings*, pages 48–52. Four Walls Eight Windows, New York, 1991.

[Watt, 1993] Alan Watt. *3D Computer Graphics*, chapter 13. Addison-Wesley, Reading, MA, 2nd edition, 1993.

[Wavish and Graham, 1996] Peter Wavish and Michael Graham. A situated action approach to implementing characters in computer games. *AAI*, 10, 1996.

[Webb *et al.*, 1981] Linda J. Webb, Carlo C. DiClemente, Ediwn E. Johnstone, Joyce L. Sanders, and Robin A. Perley, editors. *DSM-III Training Guide for Use with the American Psychiatric Association's Diagnostic and Statistical Manual of Mental Disorders (Third Edition)*. Brunner / Mazel, New York, 1981.

[Weizenbaum, 1965] Joseph Weizenbaum. Eliza — a computer program for the study of natural language communication between man and machine. *Communications of the Association for Computing Machinery*, (1):36–45, 1965.

[Wernecke, 1994] Josie Wernecke. *The Inventor Mentor: Programming Object-Oriented 3D Graphics with Open Inventor (TM), Release 2*. Addison-Wesley, Reading, MA, 1994.

[Winograd and Flores, 1986] Terry Winograd and Carlos F. Flores. *Understanding Computers and Cognition: A New Foundation for Design*. Ablex Pub. Corp., Norwood, NJ, 1986.

[Winograd, 1972] Terry Winograd. *Understanding Natural Language*. Academic Press, New York, 1972.

[Wise, 1996] J. MacGregor Wise. Intelligent agency. Society for Literature and Science, 1996.