

---

# An Attempt to Generalize AI

## Part 1 1: Explaining Dreaming

By Paul Almond

---

30 May 2010

Website:

<http://www.paul-almond.com>

E-mail:

[info@paul-almond.com](mailto:info@paul-almond.com)

This is the eleventh in a series of articles attempting an overview of how minds may work and how similar systems could be implemented in computers. Previous articles described a probabilistic hierarchy based on *patterns*. A pattern has a specification describing a set, or population, of *pattern instances*, distributed throughout a hierarchy containing the pattern instances of all the patterns. Each pattern's set of pattern instances is used to obtain statistical information for probabilistic predictions. Each pattern's population of pattern instances is to be described in a very general way, to provide a very general ontology. If the proposed view of the mind is generally correct, phenomena in human brains should make sense within this context. This article discusses dreaming within the context of the hierarchy of pattern instances. It is suggested that dreaming occurs in a partially functioning hierarchy. Functioning pattern instances that would normally receive pattern inputs from non-functional ones instead receive noise. The wide scope of the hierarchy, in terms of explaining the "self" as part of the model, similarly gives any explanation of dreaming involving the hierarchy being compromised a wide scope, economically explaining the general effects of dreaming on perception of reality, memory and behavior. If this is correct, dreaming is not profoundly different from waking experience, but only differs by a matter of degree. Looking at this another way, waking experience could be considered to be a particularly detailed dream that is strongly influenced by reality, with the "self" being "dreamed" by the hierarchical model with everything else.

# Table of Contents

1 Introduction .....	4
2 Explaining Dreaming .....	6
3 Explanation of the Extreme Scope of Dreams .....	7
3.1 The Extreme Scope of Dreaming .....	7
3.2 Acceptance of Reality in Dreams .....	8
4 Dreaming Experience as a Special Case of Waking Experience .....	9
5 The Activation Synthesis Theory .....	10
6 Conclusion .....	11
7 Bibliography .....	12

## List of Abbreviations

- AI artificial intelligence
- BERP basic, exploratory relevance process
- ERP exploratory relevance process

# 1 Introduction

This article is the eleventh in a series about artificial intelligence (AI) and how our own minds might work. The first article, *An Attempt to Generalize AI - Part 1: The Modeling System*, is available at <http://www.paul-almond.com/AI01.pdf>.<sup>1</sup> The second article, *An Attempt to Generalize AI - Part 2: Planning and Actions*, is at <http://www.paul-almond.com/AI02.pdf>.<sup>2</sup> The third article, *An Attempt to Generalize AI - Part 3: Forgetting*, is at <http://www.paul-almond.com/AI03.pdf>.<sup>3</sup>

These three articles described a hierarchy based on *patterns*, which are sets of *pattern instances*, and were intended to give an idea of how humans may model the world, plan actions and discard information from the model when it is no longer useful. The fourth article, *An Attempt to Generalize AI - Part 4: Modeling Efficiency*, which is at <http://www.paul-almond.com/AI04.pdf>, suggested that pattern instances should be allowed to have *incompletely specified pattern inputs*, so that it would be practical for the hierarchy to be “pruned” by some process seeking to maximize its relevance.<sup>4</sup> This required a *completely* probabilistic hierarchy, an issue dealt with in the fifth article of this series, *An Attempt to Generalize AI - Part 5: A Completely Probabilistic Hierarchy*, which is at <http://www.paul-almond.com/AI05.pdf>.<sup>5</sup>

That made a process to provide relevance in the hierarchy feasible. The sixth article, *An Attempt to Generalize AI – Part 6: Measuring Relevance*, which is at <http://www.paul-almond.com/AI06.pdf>, described a back-propagation process for measuring relevance in the hierarchy.<sup>6</sup> The problem is made tractable by the way in which the *action selection process*, described in the second article, *An Attempt to Generalize AI – Part 2: Planning and Actions*, works. The seventh article, *An Attempt to Generalize AI – Part 7: A Basic, Exploratory Relevance Process*, which is at <http://www.paul-almond.com/AI07.pdf>, described the *basic exploratory relevance process* (BERP), which uses this measuring process to direct the growth and pruning of the hierarchy.<sup>7</sup>

---

<sup>1</sup> Almond, P., 2010. *An Attempt to Generalize AI - Part 1: The Modeling System*. [Online] paul-almond.com. <http://www.paul-almond.com/AI01.pdf> or <http://www.paul-almond.com/AI01.doc>.

<sup>2</sup> Almond, P., 2010. *An Attempt to Generalize AI - Part 2: Planning and Actions*. [Online] paul-almond.com. <http://www.paul-almond.com/AI02.pdf> or <http://www.paul-almond.com/AI02.doc>.

<sup>3</sup> Almond, P., 2010. *An Attempt to Generalize AI - Part 3: Forgetting*. [Online] paul-almond.com. <http://www.paul-almond.com/AI03.pdf> or <http://www.paul-almond.com/AI03.doc>.

<sup>4</sup> Almond, P., 2010. *An Attempt to Generalize AI - Part 4: Modeling Efficiency*. [Online] paul-almond.com. <http://www.paul-almond.com/AI04.pdf> or <http://www.paul-almond.com/AI04.doc>.

<sup>5</sup> Almond, P., 2010. *An Attempt to Generalize AI - Part 5: A Completely Probabilistic Hierarchy*. [Online] paul-almond.com. <http://www.paul-almond.com/AI05.pdf> or <http://www.paul-almond.com/AI05.doc>.

<sup>6</sup> Almond, P., 2010. *An Attempt to Generalize AI - Part 6: Measuring Relevance*. [Online] paul-almond.com. <http://www.paul-almond.com/AI06.pdf> or <http://www.paul-almond.com/AI06.doc>.

<sup>7</sup> Almond, P., 2010. *An Attempt to Generalize AI - Part 7: A Basic, Exploratory Relevance Process*.

The eighth article, *An Attempt to Generalize AI – Part 8: Forgetting as Part of the Exploratory Relevance Process*, which is at <http://www.paul-almond.com/AI08.pdf>, removed the need for the forgetting process in the third article, instead incorporating forgetting into the BERP, or any other exploratory relevance process (ERP).<sup>8</sup> This was done by modifying the relevance measurement process (RMP) to take account of obsolescence. This article also introduced *ghost pattern instances*. A ghost pattern instance is one that persists temporarily, after “removal” by the ERP, as a simple probability value, while it is still needed as a pattern input by other pattern instances. The incompletely specified pattern inputs introduced earlier are now ghost pattern instances. The ninth article, *An Attempt to Generalize AI – Part 9: Improving the Exploratory Relevance Process*, which is at <http://www.paul-almond.com/AI09.pdf>, discussed ways in which the sophistication of the BERP might be increased, giving an improved ERP.<sup>9</sup>

Functioning of the system requires pattern instances to be placed in the hierarchy on an ongoing basis, and pattern instances need to belong to patterns. A way in which this could work had been described in the first article, *An Attempt to Generalize AI - Part 1: The Modeling System*.<sup>10</sup> Other approaches, conforming to the same general idea, were discussed in *An Attempt to Generalize AI – Part 10: Alternatives for Pattern Instance Construction*, which is at <http://www.paul-almond.com/AI10.pdf>.<sup>11</sup>

This series of articles, as well as being a development of an AI proposal, is attempting to describe how humans think. If the proposed view of the mind as a hierarchy of pattern instances is correct, or close to being correct, things known to happen with human brains should make sense within the context that it provides. This article will be looking at one of these things, dreaming, within the context of the hierarchical system described in this series.

---

[Online] paul-almond.com. <http://www.paul-almond.com/AI07.pdf> or <http://www.paul-almond.com/AI07.doc>.

<sup>8</sup> Almond, P., 2010. *An Attempt to Generalize AI – Part 8: Forgetting as Part of the Exploratory Relevance Process*. [Online] paul-almond.com. <http://www.paul-almond.com/AI08.pdf> or <http://www.paul-almond.com/AI08.doc>.

<sup>9</sup> Almond, P., 2010. *An Attempt to Generalize AI - Part 9: Improving the Exploratory Relevance Process*. [Online] paul-almond.com. <http://www.paul-almond.com/AI09.pdf> or <http://www.paul-almond.com/AI09.doc>.

<sup>10</sup> Almond, P., 2010. *An Attempt to Generalize AI - Part 1: The Modeling System*. [Online] paul-almond.com. <http://www.paul-almond.com/AI01.pdf> or <http://www.paul-almond.com/AI01.doc>.

<sup>11</sup> Almond, P., 2010. *An Attempt to Generalize AI - Part 10: Alternatives for Pattern Instance Construction*. [Online] paul-almond.com. <http://www.paul-almond.com/AI10.pdf> or <http://www.paul-almond.com/AI10.doc>.

## 2 Explaining Dreaming

With the view of cognition taken in this series, the obvious explanation of dreams involves viewing them as the results of a partially functional hierarchy. I am not claiming to be able to describe *exactly* how dreaming occurs, but the general idea is that the functionality of the hierarchy is reduced. One way in which this could work is as follows.

Suppose some pattern instances in the hierarchy are functioning properly, but many are not. Probabilistic propagation can still occur through the hierarchy, but the non-functionality of many pattern instances compromises this process. When a pattern instance's probability<sup>12</sup> would normally be affected by the probability of the pattern instance that is now non-functioning, the non-functioning pattern instance can be regarded as having a random probability: Probabilistic information propagated to functioning pattern instances from non-functioning pattern instances is just noise. This does not mean that the non-functioning pattern instance must take an active part in this: However this works, the non-functionality means that any pattern instance using it as a pattern input receives no useful information.

The noise in the hierarchy will compromise the resultant model's accuracy. If the hierarchy's functionality is extremely reduced, the model may become so inaccurate that it starts to describe a situation which is very different to reality, one which is merely a caricature of it, and this could be what dreaming is.

The hierarchy's functionality should obviously be expected to be reduced at the bottom level, where the pattern instances corresponding to external inputs/output are, but it could also be reduced at higher levels.

---

<sup>12</sup> It is being assumed, again, that pattern instances in the actual hierarchy are described by single probability values, as described in Almond, P., 2010. *An Attempt to Generalize AI - Part 10: Alternatives for Pattern Instance Construction*. [Online] paul-almond.com. <http://www.paul-almond.com/AI0.pdf> or <http://www.paul-almond.com/AI10.doc>. p.6.

## 3 Explanation of the Extreme Scope of Dreams

### 3.1 The Extreme Scope of Dreaming

An interesting aspect of dreaming is its extreme scope. I will try to explain what I find significant here.

We do not just dream narratives as passive observers. We think we are involved in the events: We are characters in our own dreams, with our own actions being dreamed. If we thought that there was some kind of external “self” or external planning system, interacting with the model, we might explain this by saying that dreaming involves a distorted model of current events, and that our “self” responds to it in the usual way, as best it can. We might imagine that dreaming involves sensory inputs being “faked” for example, so that a dream is essentially like substituting your experience for a virtual reality one, with you, the “self” continuing to operate fairly normally while interacting with this fake reality.

Dreaming, however, can go further than this. Sometimes, we might not just dream that current events are different to those in reality, but that events in the past leading up to the dream have been radically different. You might dream, for example, that you have been working in a different job for months, or that some events that, in reality happened a long time ago, are happening now. This makes it less plausible to try to explain dreams just in terms of faked sensory input. Rather, it seems that the entire model, including the representation of the past, has been compromised. This makes sense in the context of the hierarchical system described in this series of articles: The same hierarchy is used to model everything – past, present and future – so anything that compromises the hierarchy is liable to interfere with the entire description of reality.

Dreaming can have a still wider scope. We may have dreams in which not just the “outside world” is distorted, but also our own actions, so that we behave “out of character”, acting in a way in which we would never act while awake. This should be significant to anyone thinking about the relationship between the “self” and the rest of the mind: Our behavior can be distorted in dreams just as if we were part of the reality being dreamed – as if we were a character in the dream. If we think that there is a “self”, separate to the model of reality, which uses that model to plan its actions, we might explain this by saying that the “self” happens to be compromised in the same way as the model, so that there is a distorted model and a distorted self, but I suggest that this is starting to get implausible: It makes more sense to think that the “self” is just part of the model. As discussed in this series of articles, there is no processing outside the model, or indeed any special processing at all, to provide a “self”. Instead, the “self” is just the part of the hierarchical model that happens to relate past inputs/outputs to

future outputs particularly strongly. When the self is just part of the model it is natural that dreams can affect our behavior in the same way that they can affect other things. (The idea that the “self” is part of the model has been proposed previously Thomas Metzinger.<sup>13</sup>)

This does not mean that dreams are merely random noise. They can still involve propagation of probabilistic information that already exists in the hierarchy, although the extent to which the dream provides a “sensible” continuation of the past will depend on the extent to which the hierarchy is compromised.

### 3.2 Acceptance of Reality in Dreams

I just said that dreams can involve behavior that is out of character. A common example of this is the way in which we tend uncritically to accept what is going on in a dream as normal. Many dreams, of course, may be fairly normal, but many involve things that would cause us extreme surprise or concern in waking life. For example, we might have dreams in which clock faces appear distorted, or in which strange creatures appear.

Many people think that our “critical thinking” is somehow compromised in dreams, so that we are more accepting of what happens. The explanation of dreaming given here deals with this differently. The hierarchy will tend to produce the simplest model that it can find and the reason that you may accept things in a dream without criticism is that the “you” that is supposed to be doing the criticism is just being generated in the compromised model along with everything else. When the dream involves some distortion of reality we should expect other elements of the dream to be consistent with this distortion. If reality in a dream is distorted in some way, then it is actually simpler to have a dream in which characters in that dream interact with that reality as if it is normal – and you are just one of these characters.

Of course, there may be dreams in which strange things happen and characters notice this, or in which you notice this: This seems to be connected with lucid dreaming. However, expecting dreams routinely to involve characters, including you, reacting to the strangeness in the dream ignores the idea that this whole narrative includes you, because that is how you normally exist. It would be like reading a science fiction or fantasy novel set in some strange world and expecting the characters to react to this world as we would: The difficulty we have in breaking the fourth wall in dreams is strong evidence of our status as part of the model.

---

<sup>13</sup> Metzinger, T., 2003. *Being No One: The Self-Model Theory of Subjectivity*. Cambridge (MA): MIT Press.  
Metzinger, T., 2009. *The EGO Tunnel: The Science of the Mind and the Myth of the Self*. New York: Basic Books.

## 4 Dreaming Experience as a Special Case of Waking Experience

A view like this suggests that dreams are not *profoundly* different to waking experience. I am not saying that dreams are almost the same as waking experience. They are very different, but only in a matter of degree. Waking experience involves a highly functional hierarchy and dreaming involves a hierarchy with minimal functionality. In between there can be, in principle, a continuum of different states. We might imagine the hierarchy for waking experience and gradually reducing its functionality. As the number of active pattern instances decreases, the model resembles reality less and less, eventually becoming nothing more than a vague caricature of reality.

Another way of looking at this is to say that waking experience is effectively a dream, but one that is very strongly affected by the outside world. The “self” in the hierarchy would merely be part of this dream, so that when you are awake your brain is essentially “dreaming” you.

## 5 The Activation Synthesis Theory

The explanation of dreaming given here has some similarity with the *activation synthesis theory* of John Allan Hobson and Robert McCarley, which suggests that dreams are caused by random firing of neurons in the cerebral cortex and the brain making a narrative from the nonsensical sensory inputs.<sup>14</sup> Differences could be as follows.

- With what is suggested here, the hierarchy provides a specific context.
- What is suggested here is more extreme. A dream is not just a narrative made from random sensory information: The pattern instances on the bottom level of the hierarchy also provide information about the system's outputs, so the effects of this can be wider in scope.
- What we think of as the "self" is put together in the model from the bottom-level pattern instances representing previous inputs/outputs. Anything which affects the hierarchy can affect what produces *us*, making the possible scope of dreaming as wide as it could be.<sup>15</sup>
- The functionality of the hierarchy is not necessarily reduced just on the bottom level of the hierarchy.

I am not suggesting that the activation synthesis theory is wrong. Rather, I am suggesting that the kind of system discussed in these articles gives a context in which this kind of idea can have considerable explanatory power.

---

<sup>14</sup> Hobson, J.A., & McCarley, R., 1977. The brain as a dream state generator: An activation-synthesis hypothesis of the dream process. *American Journal of Psychiatry*, 134, pp.1335-1348.

<sup>15</sup> It could be argued that this is implied by the previous point anyway.

## 6 Conclusion

If the proposed view of the mind is generally correct, things known to happen with human brains should make sense within this context. One of these things, dreaming, has been considered.

It has been suggested that dreaming results from a partially functioning hierarchy, in which pattern inputs from non-functioning pattern instances are received as noise. This has some similarity with the activation synthesis theory of Hobson and McCarley.<sup>16</sup>

Dreaming within the context of the hierarchy of pattern instances makes sense, because the wide scope of the hierarchy, in terms of the “self” being a part of the model, makes anything that compromises the hierarchy similarly wide in scope. A view like this can explain why dreams can be of situations in which the past has been altered, or in which the dreamer’s (dreamed) behavior is out of character.

One way in which our behavior can be out of character in dreams is in our uncritical acceptance of what is happening in dreams. This makes sense when we view the “self” as a part of the model, rather than separate from it, so that any distortion of reality in the dream will tend to distort it as well.

This could be seen as supporting the general view of cognition in this series of articles.

---

<sup>16</sup> Hobson, J.A., & McCarley, R., 1977. The brain as a dream state generator: An activation-synthesis hypothesis of the dream process. *American Journal of Psychiatry*, 134, pp.1335-1348.

## 7 Bibliography

- Almond, P., 2010. *An Attempt to Generalize AI - Part 1: The Modeling System*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/AI01.pdf> or <http://www.paul-almond.com/AI01.doc> [Accessed 30 May 2010].
- Almond, P., 2010. *An Attempt to Generalize AI - Part 2: Planning and Actions*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/AI02.pdf> or <http://www.paul-almond.com/AI02.doc> [Accessed 30 May 2010].
- Almond, P., 2010. *An Attempt to Generalize AI - Part 3: Forgetting*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/AI03.pdf> or <http://www.paul-almond.com/AI03.doc> [Accessed 30 May 2010].
- Almond, P., 2010. *An Attempt to Generalize AI - Part 4: Modeling Efficiency*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/AI04.pdf> or <http://www.paul-almond.com/AI04.doc> [Accessed 30 May 2010].
- Almond, P., 2010. *An Attempt to Generalize AI - Part 5: A Completely Probabilistic Hierarchy*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/AI05.pdf> or <http://www.paul-almond.com/AI05.doc> [Accessed 30 May 2010].
- Almond, P., 2010. *An Attempt to Generalize AI - Part 6: Measuring Relevance*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/AI06.pdf> or <http://www.paul-almond.com/AI06.doc> [Accessed 30 May 2010].
- Almond, P., 2010. *An Attempt to Generalize AI - Part 7: A Basic, Exploratory Relevance Process*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/AI07.pdf> or <http://www.paul-almond.com/AI07.doc> [Accessed 30 May 2010].
- Almond, P., 2010. *An Attempt to Generalize AI - Part 8: Forgetting as Part of the Exploratory Relevance Process*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/AI08.pdf> or <http://www.paul-almond.com/AI08.doc> [Accessed 30 May 2010].
- Almond, P., 2010. *An Attempt to Generalize AI - Part 9: Improving the Exploratory Relevance Process*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/AI09.pdf> or <http://www.paul-almond.com/AI09.doc> [Accessed 30 May 2010].
- Almond, P., 2010. *An Attempt to Generalize AI - Part 10: Alternatives for Pattern Instance Construction*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/AI10.pdf> or <http://www.paul-almond.com/AI10.doc> [Accessed 30 May 2010].

Ibid. p.6.

## An Attempt to Generalize AI – Part 11: Explaining Dreaming

Hobson, J.A., & McCarley, R., 1977. The brain as a dream state generator: An activation-synthesis hypothesis of the dream process. *American Journal of Psychiatry*, 134, pp.1335-1348.

Metzinger, T., 2003. *Being No One: The Self-Model Theory of Subjectivity*. Cambridge (MA): MIT Press.

Metzinger, T., 2009. *The EGO Tunnel: The Science of the Mind and the Myth of the Self*. New York: Basic Books.