A Channel Theoretic Model of Situated Default Reasoning
(extended abstract)

Lawrence Cavedon

Department of Computer Science
Royal Melbourne Institute of Technology
GPO Box 2476V
Melbourne 3001
Australia

Email: `cavedon@cs.rmit.edu.au`
Fax: (+61 3) 9662-1617

December 29, 1995

A Channel Theoretic Model of Situated Default Reasoning[1]

LAWRENCE CAVEDON
Dept. of Computer Science, R.M.I.T.
LaTrobe St, Melbourne Australia
`cavedon@cs.rmit.edu.au`

1 Introduction

Defeasible reasoning is a central topic in Artificial Intelligence (AI), and has also received much recent attention in philosophy and formal semantics. Typically, this process involves reasoning with *incomplete information*—an agent leaps to some conclusion based on the information she has, whereas a different conclusion would be reached if more information about her situation was known. Procedures for default reasoning often involve some form of "internal consistency check"—the potential conclusion is accepted only if it is consistent with the information the agent has available to her. For example, if the agent knows that Tweety is a bird, then (in the absence of information to the contrary) the agent will conclude that Tweety can fly; if on the other hand she knows additionally that Tweety is a penguin, then she will instead be led to the conclusion that Tweety *cannot* fly.

In this abstract, we outline a process whereby such defeasible reasoning is grounded in *situated* regularities. Such a regularity is one which is not applicable (i.e. cannot be relied upon) in all situations, but rather is only reliable in certain circumstances. The defeasible inference process involves an agent choosing an appropriate regularity as carefully as possible (given certain information) and drawing conclusions based on this. For example, when reasoning about a general bird, an agent may choose to use the regularity "Birds fly"; however, when reasoning about a penguin, then a more appropriate regularity is "Birds which aren't penguins fly".[2]

In order to model this process, we use Barwise and Seligman's [2, 14] (henceforth B&S) *channel theory,* a recent mathematical theory of information flow that models many important features of information, such as the fallibility and situation-dependence of natural regularities. These natural regularities underly information-flow and, consequently, inference. Previous work [6, 7] has investigated the use of channel theory as a semantic model for generic sentences, which are clearly related to the default reasoning process. Further, the particular model described below is founded on a channel theoretic model of conditional logic [4]—this model involves modifying B&S's operations on channels so as to account for a regularity's "implicit background assumptions" in such a way that does not require such assumptions to be represented in the channel supporting that regularity.

This new perspective on defeasible reasoning seems to offer the possibility of significant computational advantages. The consistency checking is the source of major computational com-

---

[1] Much of the work described here was performed during the author's time as a student in the Centre for Cognitive Science, University of Edinburgh and has greatly benefited from discussions with Robin Cooper.

[2] There may be other even more appropriate regularities. This particular choice reflects the specific behaviour of the model described below.

plexity in standard default reasoning methods. The methodology outlined below allows the agent to reason *monotonically* using approximate rules. The problem then shifts to choosing the most appropriate version of a rule (i.e. based on the situation the agent finds itself in). The computational issues are outlined below—such issues are discussed further in the full paper, as is an application of the methodology to AI's *qualification problem* (i.e. the problem of specifying implicit assumptions behind rules for reasoning about the effects of actions).[3]

2 Channel Theory

In this section, I briefly review the basics of *channel theory,* B&S's theory of information flow; see [2, 14] for further details. Channel theory is concerned with *natural regularities,* informational relationships between entities and types which give structure to the world.[4] A somewhat problematic property for such regularities is that they can be both "reliable" (i.e. useful for inference) yet fallible (i.e. may admit exceptions). Such properties are, of course, required of the semantic objects used to model generic sentences. B&S resolve the problem by ensuring that relationships between the tokens of the theory play an important role—i.e. it is not just the regularity between types that counts, but also the connections between individuals that are classified by those types, and the way in which these connections are classified as instances of the general regularities.

These (fallible) regularities underly inference performed by an agent. However, by shifting the fallibility onto the regularities themselves, we obtain a different perspective on default reasoning—i.e. each of the default regularities are *bona fide* regularities, even though they allow exceptions. (This is an important consideration when discussing the semantic content of generic sentences.[5]) This differs from the usual view of default reasoning, by which it is an *agent* that is fallible—the agent only uses approximations of "true" inference rules, based on her information. The channel theoretic view more closely conforms to a notion of situated cognition, whereby the agent makes use of regularities that best suit her embedding circumstances.

2.1 Classifications

A classification is a structure that carves up (part of) the world into *tokens* and assigns various *types* to these tokens.

Definition *A* classification *A is a structure* $\langle tok(A), typ(A), :^+_A, :^-_A\rangle$ *consisting of a set of* tokens $tok(A)$, *a set of* types $typ(A)$ *and relations* $:^+_A, :^-_A$ *on* $tok(A) \times typ(A)$. *For* $t \in tok(A)$ *and* $\phi \in typ(A)$, *we say* t *is* classified positively (resp., negatively) by ϕ in A *if* $t :^+_A \phi$ *(resp.,* $t :^-_A \phi$*) holds.*[6]

The tokens of a classification may be objects, individuals, situations, or even other classifica-

---

[3] Bouquet and Giunchiglia [3] have formulated a related approach to the qualification problem.

[4] Channel theory evolved from an attempt to rectify problems with the situation theoretic account of *conditional constraints.*

[5] The channel theoretic view of generics provides other benefits, as discussed by Cavedon and Glasbey [6, 7].

[6] I will usually drop the subscripts A when this causes no confusion.

tions, while the types are properties appropriate for classifying (along some dimension) these tokens. A classification can be seen as a representation of information of a particular sort. This representation is highly relativistic—intuitively, one could view a classification as being relativised to a specific agent, epistemic state, point of view, or perspective on the world.

I assume throughout that the set of types of a classification A has internal structure: in particular, $typ(A)$ is closed under type-conjunction, type-disjunction, and type-negation operations, denoted $\wedge, \vee, \neg$ respectively. I assume these operations to satisfy certain standard properties. Since for each $\phi \in typ(A)$ there is a corresponding $\neg\phi \in typ(A)$ satisfying $t :^{+} \neg\phi$ iff $t :^{-} \phi$, I write $t :^{+} \phi$ simply as $t : \phi$ and $t :^{-} \phi$ as $t : \neg\phi$. The types are ordered under a *type-entailment* relation $\leq_A$, satisfying some obvious properties with respect to the type operations. Type-entailment constrains the classification relation in the following way: if $t : \phi$ and $\psi \leq_A \phi$ then $t : \psi$. There is also a *type-conflict* relation $\perp_A$ on $typ(A)$, involving the expected associated interaction with the classification relation and with type-entailment (e.g.: if $\phi \perp_A \psi$ then not both $a : \phi$ and $a : \psi$ hold; $\phi \perp_A \neg\phi$).[7]

2.2 Channels

A channel $\mathcal{C}$ is basically a link between two classifications A and B, licensing the flow of information of some particular sort. Formally, a channel is itself a classification—the tokens of $\mathcal{C}$ are *connections* (each linking tokens of A and B).[8]

Definition *Let A and B be classifications. A channel $\mathcal{C} : A \Rightarrow B$ linking A and B is a classification $\langle tok(\mathcal{C}), typ(\mathcal{C}), :\rangle$.*[9] *The tokens of $\mathcal{C}$ are* connections, *denoted $s \mapsto s'$, with $s \in tok(A)$ and $s' \in tok(B)$. The types of $\mathcal{C}$ are* constraints, *denoted $\phi \rightarrow \psi$, with $\phi \in typ(A)$ and $\psi \in typ(B)$.*
I sometimes refer to $\mapsto$ and $\rightarrow$ as signalling *and* indicating *relations, respectively. Given a constraint $\gamma = \phi \rightarrow \psi$, $ante(\gamma)$ denotes ϕ and $succ(\gamma)$ denotes ψ.*

Intuitively, a channel $\mathcal{C} : A \Rightarrow B$ regulates the flow of information between A and B. The type-token distinction allows a characterisation of different sorts of "erroneous" information flow.

Definition *Let $\mathcal{C} : A \Rightarrow B$ be a channel and s a token such that $(s : \phi)$ holds in A. s is a* pseudo-signal *if there is no connection $(s \mapsto s') \in tok(\mathcal{C})$. Alternatively, if there* is *a connection $(s \mapsto s') \in tok(\mathcal{C})$ and a constraint $(\phi \rightarrow \psi) \in typ(\mathcal{C})$ such that $(s \mapsto s' : \phi \rightarrow \psi)$ does* not *hold in $\mathcal{C}$, then $s \mapsto s'$ is said to be an* exception *in $\mathcal{C}$.*

Pseudo-signals and exceptions are two different ways in which we can model the failure of a channel to support "expected" information flow. These concepts are discussed in more detail

---

[7] All these concepts are more fully defined in the full paper. See also Seligman and Barwise [14].

[8] Seligman and Barwise define a channel as consisting of a classification plus two functions that map the tokens and types to their "endpoints", i.e. tokens and types of the linked classifications. In general, one should not identify the tokens and types of a channel with the pairs of tokens and types from the linked classifications, as I do here—for example, there may be two distinct connections involving the same tokens. This problem is amplified in the presence of serial composition, below. However, this abuse leads to significant simplification and abbreviation of the presentation. The full paper treats this matter more rigorously.

[9] There should of course be two classification relations here; however, I do not make use of the negative classification relation here, so I will ignore it.

in the full paper.

2.3 Operations on channels

There are various operations that can be applied to channels, resulting in more complex channels that supports more complex information flow. Given channels $\mathcal{C}_1 : A \Rightarrow B$ and $\mathcal{C}_2 : B \Rightarrow C$, we can *serially compose* them to obtain a channel that classifies information flow from A to C.[10]

Definition *Given connections* $s \mapsto s'$, $s' \mapsto s''$ *and constraints* $\phi \rightarrow \psi$, $\psi \rightarrow \tau$, *let* $(s \mapsto s' ; s' \mapsto s'')$ *be the connection* $s \mapsto s''$ *and* $(\phi \rightarrow \psi ; \psi \rightarrow \tau)$ *be the constraint* $\phi \rightarrow \tau$. *The (standard)* serial composition *of channels* $\mathcal{C}_1 : A \Rightarrow B$ *and* $\mathcal{C}_2 : B \Rightarrow C$ *is the channel* $(\mathcal{C}_1 ; \mathcal{C}_2) : A \Rightarrow C$ *such that*

1. $tok(\mathcal{C}_1 ; \mathcal{C}_2) = \{(s \mapsto s' ; s' \mapsto s'') \mid s \mapsto s' \in tok(\mathcal{C}_1) \text{ and } s' \mapsto s'' \in tok(\mathcal{C}_2)\}$;
2. $typ(\mathcal{C}_1 ; \mathcal{C}_2) = \{(\phi \rightarrow \psi ; \psi \rightarrow \tau) \mid \phi \rightarrow \psi \in typ(\mathcal{C}_1) \text{ and } \psi \rightarrow \tau \in typ(\mathcal{C}_2)\}$;
3. *the classification relation of* $(\mathcal{C}_1 ; \mathcal{C}_2)$ *is the smallest relation such that*
   $(s \mapsto s' ; s' \mapsto s'') : (\phi \rightarrow \psi ; \psi \rightarrow \tau)$ *holds in* $(\mathcal{C}_1 ; \mathcal{C}_2)$ *if*
   $s \mapsto s' : \phi \rightarrow \psi$ *holds in* $\mathcal{C}_1$ *and* $s' \mapsto s'' : \psi \rightarrow \tau$ *hold in* $\mathcal{C}_2$.[11]

Serial composition is the only channel operation defined in this abstract. The full paper considers other operations, namely *parallel composition* and *contraposition* (see [1]). These operations lead to a calculus which can be used to model complex patterns of reasoning. (This is expanded upon in the full paper.)

Cavedon [4] has shown that the standard channel operations are unsuitable for some tasks. In particular, Cavedon shows that to obtain a suitable logic of conditionals, the channel operations must be modified so as to take account of the "implicit background assumptions" of a channel. Cavedon achieves this via a *subchannel* relation: channels are related by an ordering under which $\mathcal{C}_1$ is a subchannel of $\mathcal{C}_2$ if $\mathcal{C}_2$ is somehow "more reliable" than $\mathcal{C}_1$, in that it has fewer exceptions. The central definitions are the following:

Definition *Classification*[12] *A is a* subclassification *of classification B, written* $A \sqsubseteq B$, *iff* $tok(A) \subseteq tok(B)$, $typ(A) \subseteq typ(B)$, $\leq_B$ *agrees with* $\leq_A$ *when restricted to* $typ(A)$, *and for each* $a \in tok(A)$ *and* $\phi \in typ(A)$, $a : \phi$ *holds in A iff* $a : \phi$ *holds in B.*

Definition *Let* $\phi \rightarrow \psi$, $\phi' \rightarrow \psi'$ *be constraints contained in channels* $\mathcal{C} : A \Rightarrow B$ *and* $\mathcal{C}' : A' \Rightarrow B'$, *respectively. We write* $\phi \rightarrow \psi \leq_{(A',B')} \phi' \rightarrow \psi'$, wrt $\mathcal{C}$ and $\mathcal{C}'$,[13] *if:*

1. $A \sqsubseteq A'$, $B \sqsubseteq B'$,
2. $\phi \leq_{A'} \phi'$, *and*
3. $\psi' \leq_{B'} \psi$.

---

[10] Serial composition is the channel-theoretic concept supporting Dretske's *Xerox Principle* [8].

[11] This last case can be stated a lot more succinctly in the absence of the aforementioned conflation between a connection and the pair of tokens it connects.

[12] Recall that each classification A has a type-entailment ordering $\leq_A$ associated with its types—$\phi \leq_A \psi$ if ψ entails ϕ.

[13] I drop the reference to the specific channels when this causes no confusion.

Definition *Let* $\mathcal{C} : A \Rightarrow B$ *and* $\mathcal{C}' : A' \Rightarrow B'$ *be channels, and* f *a function from* $typ(\mathcal{C})$ *to* $typ(\mathcal{C}')$. *We say* $\mathcal{C}$ *is a* $\leq_{(A',B')}$-subchannel *of* $\mathcal{C}'$ wrt f, *written* $\mathcal{C} \sqsubseteq_{(f,A',B')} \mathcal{C}'$, *if*

1. $A \sqsubseteq A', B \sqsubseteq B'$;
2. $tok(\mathcal{C}) \subseteq tok(\mathcal{C}')$;
3. *for all* $c \in tok(\mathcal{C})$ *and* $\gamma \in typ(\mathcal{C})$, *if* $c : \gamma$ *in* $\mathcal{C}$ *then* $c : f(\gamma)$ *in* $\mathcal{C}'$;
4. *for all* $\gamma \in typ(\mathcal{C})$, $\gamma \leq_{(A',B')} f(\gamma)$ *wrt* $\mathcal{C}$ *and* $\mathcal{C}'$.

We write $\mathcal{C} \sqsubseteq_{(A',B')} \mathcal{C}'$ *if there exists an* f *such that* $\mathcal{C} \sqsubseteq_{(f,A',B')} \mathcal{C}'$.

In the above definition, $\mathcal{C}'$ (i.e. the superchannel) is meant to be seen as a sort of "refinement" of $\mathcal{C}$, in that $\mathcal{C}'$ deals with the same sort of information flow as $\mathcal{C}$, but is "more discriminatory" (in that it deals with a richer set of types—this is modelled by the requirement that A and B are subclassifications of A' and B') and therefore more reliable. Note that the definition ensures that exception-hood is preserved downward—any connection that is an exception in $\mathcal{C}'$ *must* be an exception in $\mathcal{C}$ (so long as that connection is a token in $\mathcal{C}$).

The subchannel ordering allows the background assumptions of a channel $\mathcal{C}$ to be captured *without* representing them within $\mathcal{C}$ itself. Rather, such assumptions are represented in the relationship between $\mathcal{C}$ and other channels. Further, $\mathcal{C}$'s assumptions need not be completely represented in any single channel—they are distributed throughout a hierarchy of channels ordered by subchannel. (See [4, 5] for other properties of the subchannel ordering.) Given a subchannel hierarchy that encodes the background assumptions of the channels therein, the operations on channels can be modified so as to take these background assumptions into account. For example, serial composition is modified as follows:

Definition *Suppose we have channels* $\mathcal{C}_1 : A \Rightarrow B$ *and* $\mathcal{C}_2 : B \Rightarrow C$, *with* $\phi \rightarrow \psi \in typ(\mathcal{C}_1)$ *and* $\psi \rightarrow \tau \in typ(\mathcal{C}_2)$. *The* conditional serial composition *of* $\mathcal{C}_1$ *and* $\mathcal{C}_2$, *denoted* $(\mathcal{C}_1 \,;\, \mathcal{C}_2)$, *is the same as the standard serial composition of* $\mathcal{C}_1$ *and* $\mathcal{C}_2$ *except that* $\phi \rightarrow \tau \notin typ(\mathcal{C}_1 \,;\, \mathcal{C}_2)$ *if there exist* $\mathcal{C}'_1, \mathcal{C}'_2$ *such that*

1. $\mathcal{C}_1 \sqsubseteq_{f_1} \mathcal{C}'_1$, $\mathcal{C}_2 \sqsubseteq_{f_2} \mathcal{C}'_2$; *and*
2. $ante(f_1(\phi \rightarrow \psi)) \perp ante(f_2(\psi \rightarrow \tau))$.[14]

By ensuring that a channel hierarchy satisfies certain weak conditions, a channel theoretic logic of conditionals that satisfies many important properties and supports powerful patterns of reasoning is obtained (see [4, 5] for details).

3 Representing default rules in channel theory

Our channel theoretic model of default reasoning uses constraints to represent default rules. For example, the rule "(by default) birds fly" is represented by a constraint of the form $bird \rightarrow flies$, where $bird$ is a type that classifies birds and $flies$ is a type that classifies things that can fly. Such a constraint is supported by a channel $\mathcal{C}$, where $\mathcal{C}$ contains a set of connections at the token level. Some of these connections may be exceptions. Further,

---

[14] For simplicity, we assume here that any background assumptions are represented in antecedents of superchannels. In general, this assumption is not necessary (see [5]).

there may be some tokens which are pseudo-signals and which therefore are not involved in any connection. Pseudo-signals and exceptions in $\mathcal{C}$ represent birds which are in some way exceptional with respect to the given default rule.

What the simple channel theoretic representation of fallible rules does not impose is a sense of "maximal normality", i.e. that an individual is assumed to be a "normal" representative unless there is information to the contrary. For example, if we are told that Tweety is a bird (and told nothing else) then we may conclude by default that Tweety can fly. We need to define a suitable *maximal normality* condition that channels (arranged in a hierarchy) are assumed to satisfy; i.e. we need a condition that minimises the exceptions and pseudo-signals in a set of channels.

(In [5], it is *information links*—structures that are based on but have less internal structure than channels—that are assumed to be maximally normal (see [14] for the definition of links; they are also discussed in the full paper). In that model, it is the information links which are used to model an agent's reasoning and so it is more natural to modify the signalling relations of links rather than channels so as to match an agent's "expectations". This is discussed in the full paper, but we simplify the abstract by ignoring this issue in the current setting.)

The maximal normality condition defined below interacts with a given channel hierarchy. A particular channel hierarchy which is suitable for reasoning about generics is defined in [5]—the definition is reproduced in the appendix below. Briefly, this definition constructs a channel hierarchy from a given collection of channels (supporting constraints which represent default rules) by checking for type-conflicts in consequents and inserting the duals of the corresponding antecedents as implicit background assumptions. For example, the consequents of $bird \rightarrow flies$ and $penguin \rightarrow \neg flies$ conflict. Hence, $\neg penguin$ is inserted as a background assumption to $bird \rightarrow flies$—this is achieved by creating a channel $\mathcal{C}'$, supporting $bird \wedge \neg penguin \rightarrow flies$, and inserting $\mathcal{C}'$ into the channel hierarchy so that $\mathcal{C} \sqsubseteq \mathcal{C}'$, where $\mathcal{C}$ supports $birds \rightarrow flies$.[15] The basic hierarchy thus obtained is illustrated in Figure 1. The definition biases against inserting $\neg bird$ as a background assumption of $penguin \rightarrow \neg flies$ because a token's "penguin-ness" informationally implies its "bird-ness".[16] Using the channel operations (modified to account for background assumptions), a powerful logic of generics is obtained.

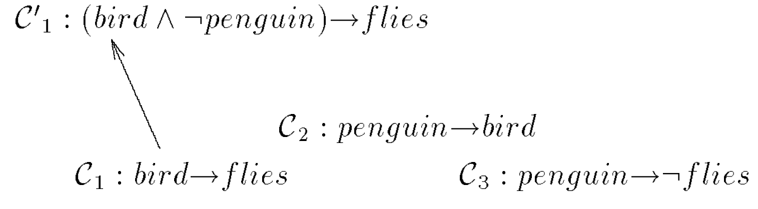

Figure 1: Channel Hierarchy for Simple Penguin Example

Having obtained a logic of generics, a logic for defeasibly reasoning about the individuals (i.e. tokens) is obtained by defining a maximal normality condition on the signalling relations

---

[15] This construction has some similarities to Pollock's use of *rebutting defeaters* [12].

[16] This is captured by the presence of a channel supporting $penguin \rightarrow bird$.

of the channels of a hierarchy[17] so as to limit the number of possible pseudo-signals and exceptions. Such a normality condition is defined with respect to a set Ψ of initially given information (i.e. a set of *Austinian* propositions of the form $(t : \phi)$). To ensure *maximal* normality, the only time that a pseudo-signal/exception should be permitted is when one of the following conditions holds:

1. the given information conflicts with the information that would otherwise be inferred;
2. inheritance from separate premises would conflict.

The second case is taken care of by referring to the channel hierarchy since conflicting consequents lead to encoded background assumptions. This leads to the following definitions. For simplicity, we assume we are dealing only with channels that relate a classification to itself (otherwise, the second case below is not needed).

Definition *Let $\mathcal{C} : A \Rightarrow A$ be a channel (which is part of a channel hierarchy) supporting a constraint $\phi \rightarrow \psi$, Ψ a collection of Austinian propositions (formed from the tokens and types of A), and $t \in tok(A)$ a token such that $(t : \phi) \in \Psi$. The connection $t \mapsto t$ is said to be* inappropriate in $\mathcal{C}$ wrt Ψ *if either:*

1. *$(t : \psi') \in \Psi$ and $\psi \perp \psi'$, or*
2. *there exists a channel $\mathcal{C}'$ such that $\mathcal{C} \sqsubseteq_f \mathcal{C}'$, where $\phi \rightarrow \psi \in typ(\mathcal{C})$ and $f(\phi \rightarrow \psi) = \phi' \rightarrow \psi'$ and $(t : \tau) \in \Psi$, where $\tau \perp \phi'$.*

To illustrate the second case, if Ψ contains $(tweety : penguin)$ then $tweety \mapsto tweety$ will be inappropriate in $\mathcal{C}_1$ of Figure 1 since this information conflicts with the antecedent of the constraint in $\mathcal{C}'_1$. On the other hand, *tweety* will be "normal" in $\mathcal{C}_3$ and $(tweety : \neg flies)$ can be inferred via this channel.

Definition *A collection $\mathcal{L}$ of channels arranged in a subchannel hierarchy is said to be* maximally normal wrt Ψ *(a collection of Austinian propositions) if for each channel $\mathcal{C} : A \Rightarrow A$ in $\mathcal{L}$ and each token $t \in tok(A)$, $t \mapsto t$ is a non-exceptional connection in $tok(\mathcal{C})$ if and only if $t \mapsto t$ is not inappropriate in $\mathcal{C}$.*[18]

Maximal normality ensures that the *only* exceptional tokens are those that are inappropriate, given a channel hierarchy and some initially given information. Notice that the definitions of inappropriateness and maximal normality are both "static"—unlike in Default Logic [13] or Commonsense Entailment [11], there is no need to define anything "iteratively". This is because the channel hierarchy is assumed to have been closed under the channel operations. This effectively partitions the default reasoning process into two stages—(i) defining the channel hierarchy and closing it under the operations, and (ii) imposing the normality condition. An advantage of this process is that (i) is only concerned with the type-level constraints (which represent default rules) and not with the tokens or the observed information at all. The second stage (i.e. definition of maximal normality) is the only part of the process that varies with different initial observed information.

Having imposed the maximal normality condition on a channel hierarchy, a channel theoretic

---

[17] As pointed out above, this should really be on the signalling relations of the corresponding information links.

[18] Actually, we will assume that there are some channels, called *logical* channels, supporting logical relationships (e.g. $\phi \wedge \psi \rightarrow \psi$) which contain *no* exceptions or pseudo-signals.

system of default inference can be obtained. (This will be defined more completely in the full paper.) This system of default inference can be shown to satisfy many important properties. For example, it satisfies the principle of Specificity (as demonstrated by the "penguin" example),[19] Graded Normality (i.e. individuals can be "abnormal" in some respects while being "normal" in others), and other important general properties of default reasoning. Cavedon [5] shows that (under a slight extension) the channel theoretic system of default reasoning effectively defines a *preferential consequence relation* in the sense of Kraus et al. [10].[20]

4 A methodology for situated default reasoning

The previous section outlined a channel theoretic model of defeasible reasoning, based on a maximal normality condition defined with respect to a channel hierarchy and some initially given information. Channel theory is intended to capture certain notions of situated cognition—in particular, the regularities which an agent relies upon to carry information about its world are reliable only within certain contexts, when certain conditions are in place.

A channel hierarchy contains (moving up the hierarchy) increasingly reliable regularities; i.e. channels which contain fewer exceptions as we move up the hierarchy. Two channels which are related by the subchannel relation are concerned with similar sorts of information flow, with the superchannel being a more reliable version.[21] However, the subchannels tend to be more efficient to use, since they deal with simpler types and a more restricted view of the world. Efficient cognitive agents are designed to use the regularities that are best suited (i.e. strike a balance between reliability and efficiency) to their environment. In the channel theoretic model, this involves choosing the level in a hierarchy which is most appropriate to the particular circumstances in which the agent finds itself.

As a simple illustration, consider again the penguin example. If an agent is reasoning about an arbitrary bird b, then the most appropriate level in the hierarchy of Figure 1 is $\mathcal{C}_1$—no information which the agent has available warrants it reasoning about the "penguin-ness" of b and therefore it should not do so. However, if the agent is reasoning about a penguin p, then it should not be led to the conclusion that p can fly. This is achieved if the appropriate level of the hierarchy is taken to be $\mathcal{C}'_1$—since the "penguin-ness" of p *is* relevant at this stage, then a channel that considers this information is the more appropriate one. Use of channel $\mathcal{C}_3$ then results in the desired conclusion $p : \neg flies$.

The process of choosing the appropriate level of the hierarchy is reflected in the definition of inappropriateness given earlier: $p \mapsto p$ is inappropriate in $\mathcal{C}_1$ but $b \mapsto b$ is not. Alternatively, the choice of appropriate channel is made by considering what information is available: this is captured by the following definition of *adequacy*.

---

[19] This principle tends to be an important benchmark for nonmonontonic logics. Systems of default reasoning which are based on a conditional logic foundation tend to satisfy this principle. Note that the described system handles *full* Specificity—some systems (e.g. that of Morreau [11]) give unintuitive results when the subtype relationship (e.g. *penguin* is a subtype of *bird*) is itself defeasible.

[20] Cavedon conjectures that the system defines a *rational* consequence relation [9] but has not yet proved this.

[21] Even though this discussion seems to suggest that superchannels are somehow more "correct" than their subchannels, it should be stressed that each channel is to be seen as supporting *bona fide* regularity. While the subchannels are less reliable, they are *not* to be seen as simply approximations to "correct" versions.

Definition *Given a collection* Ψ *of Austinian propositions, a channel* $\mathcal{C} : A \Rightarrow A$ *is* adequate *with respect to* Ψ *if for each* $(t : \phi) \in \Psi$, $t \in tok(A)$ *and* $\psi \in typ(A)$.

A channel $\mathcal{C}$ is adequate with respect to a collection Ψ of propositions if the associated classification A contains all the types and tokens that are used in the propositions in Ψ—effectively, the "language" of A is rich enough to represent all the information in Ψ. To ensure that the simplest channel is chosen, the appropriate channel for a given inference should be a minimal adequate one for the propositions in question.

The notion of adequacy of a channel in a hierarchy suggests an efficient method for defeasible reasoning which corresponds to the semantic model described in the previous section. The shift in emphasis from standard default reasoning is significant and seems to more naturally correspond to the sort of cognitive process we undertake. This is well illustrated by application of the approach to AI's *qualification problem* (the application is discussed in the full paper; see also [5]): when I attempt to start my car, I don't (usually) consider what *may* cause it to fail to start and, having no information that any such conditions hold, assume that everything is OK; I simply turn the ignition-key and expect the car to start. There is no process of explicitly reasoning about the lack of information regarding the background conditions in place; yet the usual default logic approach reflects such a process. The channel theoretic model, on the other hand, reflects a process whereby an agent uses the simple (yet fallible) regularity that suits the normal circumstances.[22]

The advantages of this methodology, from a computational point of view, is that much of the processing can be performed "off-line"—as pointed out earlier, the construction of the channel hierarchy is concerned only with the type-level constraints and does not depend on any circumstantial information. When inference is called for, the agent uses constraints which are as simple as its embedding circumstances allow. The potential complexity of the methodology involves the choice of appropriate level. However, it seems that the channel hierarchy itself can be embellished with information which makes this choice simpler. An algorithm for constructing an appropriate hierarchy and for performing defeasible reasoning (and the formal properties of it) is currently being investigated (initial results are described in the full paper).

References

[1] J. Barwise. Constraints, channels and the flow of information. In P. Aczel, D. Israel, Y. Katagiri, and S. Peters, editors, *Situation Theory and its Applications III.* CSLI Lecture Notes 37, Stanford: CSLI Publications, 1993.

[2] J. Barwise and J. Seligman. The rights and wrongs of natural regularity. In J. Tomberlin, editor, *Philosophical Perspectives, Vol. 8.* Ridgeview, California, 1994.

[3] P. Bouquet and F. Giunchiglia. Reasoning about theory adequacy: a new solution to the qualification problem. 1994. To appear in *Fundamenta Informaticae.* Also available by ftp as `ftp.mrg.dist.unige.it:pub/mrg-ftp/9406-13.ps.Z`.

[4] L. Cavedon. A Channel-Theoretic model for conditional logics (extended abstract). In *Conference on Information-Oriented Approaches to Logic, Language and Computation*, Moraga, CA, 1994. To appear in *CSLI Lecture Notes* series.

---

[22] Of course, if this regularity fails then the agent may need to revise its thinking—see [5].

[5] L. Cavedon. *A Channel Theoretic Approach to Conditional Reasoning.* PhD thesis, Centre for Cognitive Science, University of Edinburgh, 1995.

[6] L. Cavedon and S. Glasbey. Outline of an information-flow model of generics. In *Fifth Symposium on Logic and Language*, Hungary, 1994. Revised version to appear in special issue of *Acta Linguistica.*

[7] L. Cavedon and S. Glasbey. The role of context in the interpretation of generics. In *Tenth Amsterdam Colloquium on Logic and Language*, 1995.

[8] F. Dretske. *Knowledge and the Flow of Information.* Cambridge: MIT Press, 1981.

[9] M. Freund, D. Lehmann, and P. Morris. Rationality, transitivity and contraposition. *Artificial Intelligence*, 52, 1991.

[10] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1), 1990.

[11] M. Morreau. *Conditionals in Philosophy and Artificial Intelligence.* PhD thesis, Department of Philosophy, University of Amsterdam, 1992.

[12] J. L. Pollock. Defeasible reasoning. *Cognitive Science*, 11, 1987.

[13] R. Reiter. A logic for default theory. *Artificial Intelligence*, 13, 1980.

[14] J. Seligman and J. Barwise. *Channel Theory: Towards a Mathematics of Imperfect Information Flow* (draft). Available by ftp as `phil.indiana.edu:/pub/SelBar93.ps`, 1993.

A A channel hierarchy for reasoning about generics

This section contains the definition of a particular channel hierarchy that seems well-suited to reasoning about generics. Given an initial collection of channels (representing generics or default rules), a hierarchy is constructed according to the definition below. The channel operations (revised so that they take account of the hierarchy) are then applied to infer new channels (i.e. new generics).

Definition *Let $\mathcal{S}$ be a collection of channels. Let $\mathcal{T}(\mathcal{S})$ be the collection of channels obtained by closing $\mathcal{S}$ under the standard channel operations.*[23] *The* default channel hierarchy for $\mathcal{S}$ *is the smallest hierarchy $\mathcal{K}$ of channels containing $\mathcal{S}$ such that:*

> *for each channel $\mathcal{C}_1 : \phi \rightarrow \psi$ in $\mathcal{S}$ and each channel $\mathcal{C}_2 : \tau \rightarrow \psi'$ in $\mathcal{T}(\mathcal{S})$ such that $\psi \perp \psi'$, $\mathcal{K}$ contains a channel $\mathcal{C}'_1 : (\phi \wedge \neg\tau) \rightarrow \psi$, where $\mathcal{C}_1 \sqsubseteq_f \mathcal{C}'_1$ and $f(\phi \rightarrow \psi) = (\phi \wedge \neg\tau) \rightarrow \psi$;* unless *either*
>
> 1. *there is a channel $\mathcal{C} : \phi \rightarrow \tau$ in $\mathcal{T}(\mathcal{S})$;* or
> 2. *$\mathcal{C}_2$ is formed by a composition involving a channel $\mathcal{C}' : \phi \rightarrow \sigma$, for some σ.*

This choice of channel hierarchy ensures that powerful patterns of reasoning are supported by the resulting logic of generics—see [5] for an informal discussion of the intuitions behind this particular definition. Note the reasoning process does not in any way depend on the particular vagaries of the specific channel hierarchy chosen—the channel operations simply assume the existence of *some* channel hierarchy. The above definition simply happens to define a particular hierarchy that ensures powerful patterns of reasoning are satisfied.

---

[23] There is no circularity in this requirement—the *standard* channel operations do not take account of the encoded background assumptions.