A Formal Framework for Identity in Cyber and Other Universes

Albert Esterline[1] and Jules Chenou[2]

1. Department of Computer Science

2. Department of Electrical and Computer Engineering

College of Engineering

North Carolina Agricultural and technical State University

Greensboro, North Carolina, 27411

{esterlin, jchenou}@ncat.edu

Abstract

We present preliminary work on a framework for constructing identities that takes as its point of departure Barwise's channel theory, founded on category theory, and allied work on situation semantics. The framework takes actions relating to identity as fundamental. Actions and events are situations, although situations can include larger context. Identity-related situations are classified under (identity) types, which can be fused. Identity-related activity is partitioned according to the 'form of life' involved into (identity) environments, similar to Barwise's classifications. An identity-related action is often based on another identity-related action as a resource, and all these actions may reference events in an individual's life. The result is a layered picture of identity-related situations. Identities retain references to the situations from which they are constructed, so provenance, context, and narrative detail are retained. The category-theoretical constructs of Barwise's channel theory, however, must be relaxed when identity in multiple environments is considered.

1 Introduction

This paper presents preliminary work on a framework for constructing identities that takes its point of departure work by Barwise and his associates on the "flow of information" and situations. Barwise's channel theory is an application of an area of abstract mathematics known as category theory, which addresses the form of mathematical theories and interrelates their structures. Channel theory gives an account of how x being P carries the information that y is Q, where x and y are "tokens" and P and Q are "types," x and P being in one classification system, y and Q being in another. Barwise and his associates also developed situation semantics, which accounts for the meaning of an expression as the relation between situations, viz., between an utterance (situation) and a described situation. Channel theory and situation semantics were two parts of a coherent program, and situation theory was developed as mathematical support for situation semantics.

The framework sketched here takes actions relating to identity, both judgments and assertions, as fundamental. Actions are events (with an agent), and events are basically situations, although situations can include larger regions of space and time and generally accommodate context. Identity-related situations are classified under (identity) types, which can be fused. Identity-related activity is partitioned according to the 'form of life' involved into environments, similar to Barwise's classifications. An identity-related action is often based on another identity-related action as a resource, and all these actions may reference events in an individual's life. The result is a layered picture of identity-related situations, and this layering was part of the motivation for situation semantics.

Identities in this framework retain references to the situations from which they are constructed, so provenance, context, and narrative detail are retained. Multiple identities are easily accommodated yet fusion results in constraints that converge to real-life individuals behind the perspectives that arise in the environment, somewhat as a genotype is behind phenotypes.

It turns out that the category-theoretical constructs of Barwise's channel theory must be relaxed when identity in multiple environments is considered. We explain the reasons for this and sketch what needs to be done to recover the category-theoretic foundation.

The next two sections provide brief introductions to category theory and channel theory. Section 4 introduces the technical notion of a situation. Section 5 introduces the notion of identity as a type in a classification (system), that is, an environment, which involves a coherent set of

practices and standards used to identify and characterize individuals. Different environments provide different perspectives on a given person, but note that the same id-relevant situation (or id-situation) may be classified under several types since it may involve several individuals. Section 5 concludes with a sketch of three environments. Section 6 reconciles identification-relevant actions (id-actions) that are first person assertions with third-person identifications and characterizations in terms of the notion of theory of mind from developmental psychology. Section 7 introduces resource situations as meta-representations in an id-situation, and it introduces referenced events as real-world facts that are the background content of the id-situations. The next section analyzes examples where identities (as types in environments) are fused. Finally, Section 9 considers how this framework can be expressed in category-theoretic terms similar to channel theory.

2 Category Theory

A category $\mathcal{C}$ consists of a class of objects and a class of morphisms (or arrows or maps) between the objects. Each morphism *f* has a unique source object a and target object b; we write $f: a \to b$. The composition of $f: a \to b$ and $g: b \to c$ is written as $g \circ f$ and is required to be associative: if in addition $h: c \to d$, then $h \circ (g \circ f) = (h \circ g) \circ f$. It is also required that, for every object *x*, there exists a morphism $1_x: x \to x$ (the identity morphism for *x*) such that, for every morphism $f: a \to b$, we have $1_b \circ f = f = f \circ 1_a$. It follows from these properties that there is exactly one identity morphism for every object. A functor from one category to another is a structure-preserving mapping, preserving the identity and composition of morphisms. More exactly, if $\mathcal{C}$ and $\mathcal{D}$ are categories, then a functor *F* from $\mathcal{C}$ to $\mathcal{D}$ is a mapping that associates with each object $x \in \mathcal{C}$ an object, $F(x) \in \mathcal{D}$ and, with each morphism $: x \to y \in \mathcal{C}$, a morphism $F(f): F(x) \to F(y) \in \mathcal{D}$ such that $F(id_x) = id_{Fx}$ for every object $x \in C$ and $F(g \circ f) = F(g) \circ F(f)$ for all morphisms $f: x \to y$ and $g: y \to z$.

In category theory, a commutative diagram is a diagram of objects (as vertices) and morphisms (arrows between objects) such that all directed paths in the diagram with the same start and end points lead to the same result by composition.

The classic presentation of category theory is (MacLane 1998). Two reasonably comprehensive and rigorous texts that are not too sophisticated for someone with a strong undergraduate math background are (Simmons 2011) and (Awodey 2010). A light introduction is provided by (Lawvere and Schanuel 2009). (Pierce 1991) and (Barr and Wells 1990) are texts addressed specifically to computer scientists; (Fiadeiro 2005) addresses category theory in the context of software engineering.

3 Channel Theory

The point of departure for channel theory is the notion of a classification, which consists of a set of individuals ("tokens"), a set of types, and a binary relation indicating that a given token is of a given type. Barwise and Seligman (Barwise and Seligman 1997) presented a framework for the "flow of information" in (generally implicit) category-theoretic terms. They address the question, "How is it that information about any component of a system carries information about other components of the system?" They define a classification $\boldsymbol{A}$ to be a structure with non-empty sets $typ(\boldsymbol{A})$ of types and $tok(\boldsymbol{A})$ of tokens as well as a binary relation $\vDash_A$ between $tok(\boldsymbol{A})$ and $typ(\boldsymbol{A})$ such that, for $a \in tok(\boldsymbol{A})$ and $\alpha \in typ(\boldsymbol{A})$, $a \vDash_A \alpha$ indicates that *a* is of type α. The theory does not limit what *a* or α might be (as long as it makes sense for *a* to be of type α). It could be that *a* is an object and α a property (monadic first-order relation), or *a* might be a situation and α a type of situation; often, different tokens of a classification amount to the same physical system across different time points and types are instantaneous partial state descriptions of the system.

For classifications $\boldsymbol{A}$ and $\boldsymbol{C}$, an infomorphism *f* from $\boldsymbol{A}$ to $\boldsymbol{C}$ is a pair of functions $(f^\wedge, f^\vee), f^\wedge: typ(\boldsymbol{A}) \to typ(\boldsymbol{C})$, and $f^\vee: tok(\boldsymbol{C}) \to tok(\boldsymbol{A})$ satisfying, for all tokens $c \in tok(\boldsymbol{C})$ and all types $\alpha \in typ(\boldsymbol{A})$,

$f^\vee(c) \vDash_A \alpha$ iff $c \vDash_C f^\wedge(\alpha)$

That is, the image of token *c* (under $f^\vee$) is classified to be of type α if and only if *c* is classified to be of the type that is of the image of α (under $f^\vee$). Intuitively, an infomorphism is a part-to-whole, A-to-C, informational relationship. Figure 1.a shows function pairs $(f^\wedge, f^\vee)$ and $(g^\wedge, g^\vee)$, each depicted in Figure 1.b as a single infomorphism with codomain $\boldsymbol{C}$, while *f* has domain $\boldsymbol{A}$ and *g* domain $\boldsymbol{B}$.

Information (not the "amount" of information, as per Shannon) is assumed to "flow" among the components of a system. Components may, but need not, be distant from one another in time and space, and they may be very different things. The system is "distributed" in this sense (not necessarily in the sense in which that term is used in computer science). For exemplar, the students, classrooms, scheduling system, and attendance records at a school together form a distributed system.

An *information channel* is a family of infomorphisms with a common codomain, called the core. Essentially, a channel consists of a set $\{A_1,...,A_n\}$ of classifications that represent the parts of the distributed system, a classification C (the core) that represents the system as a

whole, and a set of infomorphisms $\{f_1,...,f_n\}$ from each of the parts onto C. Tokens in C are the connections of the system: a given token c in C connects the tokens it is related to by means of $\{f_1,...,f_n\}$. Parts $\{A_1,...,A_n\}$ carry information about each other as long as they all are part of C. In category-theoretic terms, the core is the *sum* (or *colimit*) of the parts. Figure 2 shows the sense in which the core is a *universal* construction, making it unique.

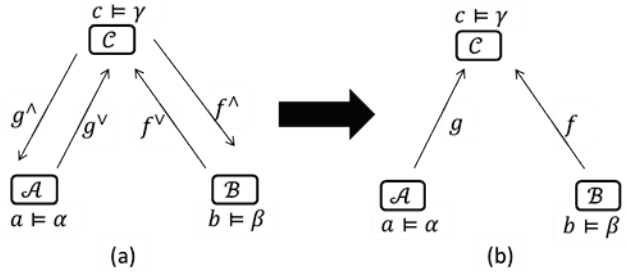

Figure 1. (a) shows infomorphisms f (with domain A) and g (domain B) with common comdain C as function pairs. (b) is the single-arrow view.

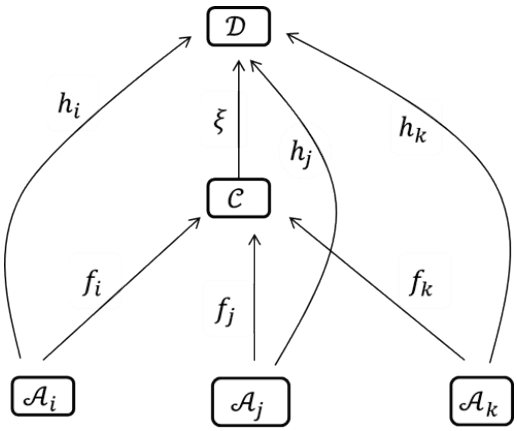

Figure 2: A core (universal cocone) $(\boldsymbol{f_i : \mathcal{A}_i \to \mathcal{C}})_{i \in I}$ *is a family of informorphisms with a common codomain* $\boldsymbol{\mathcal{C}}$*; if* $(\boldsymbol{h_i : \mathcal{A}_i \to \mathcal{D}})_{i \in I}$ *is a another such familywith a common codomain* $\boldsymbol{\mathcal{D}}$ *then there exists a unique informorphism* $\boldsymbol{\xi : \mathcal{C} \to \mathcal{D}}$ *such that for all* $\boldsymbol{i \in I, h_i = \xi \circ f_i}$.

A *distributed system* D is a collection of elements that carry information about each other. Formally, D consists of an indexed class $cla(D)$ of classifications together with a class $inf(D)$ of infomorphisms whose domains and codomains are all in $cla(D)$. An information channel K covers distributed system D if and only if $cla(D)$ are the classifications of the channel and, for every infomorphism f in $inf(D)$, there are infomorphisms from both the domain and codomain of f to the core of K such that the diagram formed by these three infomorphisms commutes. Basically, all classifications in D are "informational parts" of the core whose channel covers D.

Turning to regularities in a classification's type, let A be a classification and Γ and Δ be sets of types in A. A token *a* of A satisfies the "sequent" ⟨Γ,Δ⟩ provided that, if *a* is of every type in Γ, then it is of some type in Δ. If every token of A satisfies ⟨Γ,Δ⟩, then Γ is said to entail Δ and ⟨Γ,Δ⟩ is called a constraint supported by A. The set of all constraints supported by A is called the complete theory of A, denoted by Th(A). These constraints are system regularities, and it is by virtue of regularities among connections that information about some components of a distributed system carries information about other components. These regularities are relative to the analysis of the distributed system in terms of information channels. Barwise and Seligman's summary statement of their analysis of information flow, restricted to the simple case of a systems with two components, a and b, is as follows.

> Suppose that the token *a* is of type α. Then *a*'s being of type α carries the information that *b* is of type β, relative to channel C, if *a* and *b* are connected in C and if the translation α′ of α entails the translation β′ of β in the theory Th(C), where C is the core of C. (Barwise and Seligman 1997, p.35)

4 Situations

Channel theory is neutral about what tokens are, for example, objects (or object states) or system states. One alternative, which is quite common and natural, is to take tokens to be situations. Situation semantics, articulated in Barwise and Perry's *Situations and Attitudes* (Barwise and Perry 1984), attempted to provide a solid theoretical foundation for reasoning about commonsense and real world situations and was complemented by Barwise's other work in channel theory. Situation theory provides the mathematical foundations to situation semantics and was developed by Barwise and others, including Devlin in *Logic and Information* (Devlin 1991).

According to Devlin, a *situation* is some part of the ongoing happenings in the world and allows situation theory to handle context. Situations include not only connected space-time regions but also spatially disconnected (e.g., telephone calls) and temporally disconnected situations (e.g., multi-day sporting events).

Situation theory introduces the notion of an *infon* as the basic item of information. We assume as given sets *A* of individuals, *R* of properties and relations, *L* of spatial locations, and *T* of temporal locals. The general form of an infon, then, is

$\langle\langle R_n, a_1, \ldots, a_n, l, t, i\rangle\rangle$,

where $R_n \in R$ is an *n*-place relation, $a_1, \ldots, a_n \in A$ are objects appropriate for the corresponding argument places

of R_n, $l \in L$, $t \in T$, and $i \in \{0,1\}$ is the *polarity*. A polarity of 1 indicates that the objects are thus related in l at t; a polarity of 0 indicates otherwise. Where s is a situation and σ an infon, it might be that s *supports* σ, that is, that σ *holds* in s.

A *real situation* (according to Devlin) is a part of reality (and considered a single entity) that supports an indefinite number of infons, while an *abstract* situation is a set of infons. Real situations are tokens for us. An abstract situation for us is a type. If abstract situation s is the set $\{\sigma_1, \ldots, \sigma_m\}$ of infons and we identify the classification as $\boldsymbol{I}$, then, for real situation r, $r \models_{\boldsymbol{I}} s$ if and only if r supports all infons in s.

5 Identity as Classification

We take an identity to be a classification in the sense of channel theory, where types are abstract situations and tokens are real situations. The similarity between situations and events has been frequently noted (cf., e.g., Kratzer 2011), and actions are a kind of event (where one can identify an agent). We take situations as they relate to identity (what we call *id-situations*) to be those that include identity-relevant actions (what we call *id-actions*).

An (identity) *environment* is a coherent set of practices and standards used to identify and to characterize individuals. This notion of environment relates to Wittgenstein's observation that "to imagine a language means to imagine a form of life" (Wittgenstein 1991, §19) if we extend this observation to cover not only grammatical utterances but any action that has significance by virtue of social convention. Below, we sketch some environments by which one establishes what we call online identity, physiological identity, or social identity. Our analysis to date is modest, and the example environments are concrete. The 'forms of life' involve id-actions that are overt and subject to conventions that have evolved to sanction (sometimes legally) evidence for identity.

Different environments provide different perspectives on a given person. The question of how these perspectives are fused into a single identity is deferred to Section 8. How one partitions identity-related activity into environments is somewhat arbitrary, and our example environments could easily be partitioned finer. Indeed, we mention fusing identities within each of these environments.

The same id-situation may be classified under several types, as when an action relates to the identities of several different individuals. One generally tries to identify id-situations that relate to as few distinct identities as possible although some events inherently involve multiple players.

Some id-actions amount to asserting an identity (which is a first-person perspective) while others identify or characterize an individual distinct from the identifying or characterizing agent. We handle both these perspectives within a given environment.

5.1 Online Identity

Our first example of an identity environment is that in which we establish online identities in a rather narrow sense. Here we can look at the FOAF properties that have *Person* (or, more broadly, *Agent*) as domain. This includes, for example, *mbox* (a person's email address) and *workplaceHomepage*. As per FOAF (and the Semantic Web in general), resources are generally identified with URIs. A type is associated with a pattern that (in the first instance) is a URI. As a situation type, such a type is a context in which the URI is used (e.g., a given email address is used in the 'To' field of an email). Tokens are situations where someone engages in an online activity or reflects on such an activity. Such a situation is classified as of a given type if the activity involved makes use of the URI in question. The situation could involve an individual asserting their identity (as in the 'From' field of an email) or characterizing an individual (as in the 'To' field of an email or when someone analyzes the content of an email).

5.2 Physiological Identity

Another identity environment is that in which one establishes what we refer to as physiological identity. Types here relate to the physical features of a person that may figure in their identification or characterization, such things as fingerprints, iris patterns, photographs of faces, and recordings of voices. Thus, for example, we have a situation type in which a given fingerprint pattern is used for identification of an individual. A token here is a situation in which one uses physical evidence for identification or characterization. Perhaps the more common cases are where one identifies or characterizes an individual, but there are notable cases where one asserts one's identity (e.g., "shows one's face"). To ground the notion of a pattern and relate it to officially sanctioned evidence, we distinguish an instance of a pattern that is preserved by an authority (such as a law-enforcement agency) and accepted as a standard against which other instances of the pattern may be compared for authenticity. Examples of such standards include, for example, the images that the FBI keeps of convicted criminals' fingerprints. The pattern instantiated by such a standard may also be instantiated, for example, on a doorknob (in, say, a crime scene investigation), but it is also instantiated in the texture of the thumb of an actual person. Note that a standard need not be pictorial. For example, an iris pattern may be stored as a multitude of properties and distributions in computer memory.

5.3 Social Identity

Our final example of an identity environment is where we establish what we call (in a narrow sense) social identities, concerned with social aspects of a person that may figure

in an attempt to identify or characterize them. Paradigmatic examples come from historical research. A token here is a situation in which one uses social evidence for identifying or characterizing an individual. Types include such things as the observing a given name in a given military roster or a church record. A name is to be understood in terms of its denotation independent of language and variants within the same language. Thus, for example, 'Charlemagne', 'Carolus Magnus', and 'Karl der Große' all count as the same name. On the other hand, a phrase structured as a proper name (e.g., "John Smith") often (if not usually) denotes different individuals in different circumstances and may not denote a real individual in any circumstance (e.g., "Lieutenant Kijé"). A definite description is a way of identifying an individual based on characteristics that are distinguishing in the context in which they are used. Extended discourse can provide rich characterization of an individual, who is denoted with various linguistic devices, including proper name, definite descriptions, and anaphoric pronouns. While examples of actions relevant to social identity that come readily to mind generally involve identifying or characterizing an individual, there are notable cases of asserting social identity (e.g., as per Shelley, "My name is Ozymandias, king of kings: Look on my works, ye Mighty, and despair!").

Note that physiological identity lacks the conventional or social nature of the other two. In some standard sense of "identity," it is quite common for people to have multiple online identities, and FBI Associate Director Mark Felt was revealed to be the character Deep Throat only 31 years after Watergate. In contrast, perspectives that apparently have a life of their own do not arise with physiological evidence unless the person himself undergoes radical change.

6 Identity Environments and Theory of Mind

We have maintained that, for a given identity environment, there are id-actions that are assertions of identity and others that identify or characterize an individual distinct from the identifying agent. That the same 'form of life' is involved in both aspects (when the identity environment is held constant) is supported by research in developmental psychology addressing what is known as theory of mind (ToM) (Flavell 2004), the human ability to attribute mental states (beliefs, desires, intentions, etc.) to oneself and others and to understand that others have mental states different from one's own. One area covered by ToM research is false belief. For example, a child watches as a puppet sees a cookie put in one of two boxes and leaves. Someone moves the cookie to the other box. When the puppet then returns, the older child, but not the younger, says the puppet will look in the original box. This is because only the older child has a notion of false belief, here that the puppet may falsely believe that the cookie is in the box where it saw the cookie placed. Other areas of ToM research include perspective taking, pretend play, and recognizing the referents of words.

A quite comprehensive coverage of ToM is presented by Alvin Goldman in *Simulating Minds* (Goldman 2008). Goldman uses the term "mindreading" for representing another as having a certain mental state, and he covers the two main accounts of mindreading. According to *theory theory* (TT), one uses a naïve psychological theory to guide one's assignments of mental states, while according to *simulation theory* (ST), one tries to replicate others' mental states and project the resulting pretend states onto them. The simulation advocated by ST is not like the computational simulation of, say, a weather system, but rather is *replication* (or duplication).

Regarding identity-relevant actions, what is known of ToM suggests that, in identifying or characterizing an individual distinct from the agent doing the identifying or characterizing, the agent uses the ToM relevant to the identity environment in question. According to TT, all states of the mindreader are *meta-representations*, none of which are attributed to the target, although the *contents* of the final state are (Goldman 2008). When an agent identifies or characterizes another individual, this action generally involves a meta-representation by the agent of an identity-relevant action by that individual.

7 Resource Situations and Referenced Events

The denotations of the meta-representations invoked by TT are essentially what Devlin identifies as *resource situations*. For example, the utterance "The dog I saw yesterday is back" carries information about the presence of a dog and does so by reference to a previous situation (in which the same dog was present). Regarding id-actions, when one inspects an email, for example, and notes that it is from someone, the id-action is the inspection and the resource situation is the sending of the email. When an id-action asserts identity, no resource situation is generally needed, but when it identifies or characterizes an individual, it generally (but not always) references a resource situation. An example of the latter that does not involve a resource situation is taking a picture of an individual. Taking a picture, however, can be a resource situation for any number of id-situations because that id-action leaves an enduring object.

Another component in many id-situations is what we call the *referenced event* (or, sometimes, the referenced state), a real-world fact that is the background content of the id-situation. For example, the content of an email often

makes reference to one or more events, fingerprints on the doorknob may relate to someone opening the door at a certain time, and analyzing a passage in a chronicle usually involves reading about events that happened in the society being chronicled. Frequently, a referenced event is what a resource situation is about. It may also be what an assertion of identity is about (as in, e.g., a confession). An id-action often involves a referenced event without a resource situation as when someone describes another's personality, background, or work. When an object is provided for reference, we take the referenced 'event' to be the event producing the object. For example, when a mug shot is attached to an email, the referenced event is the subject posing for the mug shot (which includes provenance and topic information missing in the bare picture).

For establishing identity, sometimes we focus on the id-situation itself, sometime on resource situation, and sometimes on the referenced event. Times and spatial locations are often details that bind together events to form a coherent picture, and these parameters are relevant to all three of these aspects. These details, however, are often lacking for some or all of these aspects. In such cases, sometimes we can impose a non-metric (and possibly only partial) temporal order on situations and events and determine relative locations. Spatial and temporal inclusion and overlap relations are also often indicated.

8 Fusing Identity Types

Often the more interesting issues that arise regarding identity relate to fusing identities from several identity environments. Suppose we do the following.

(i) We observe an email from **`John.Doe@Acme.com`** that states that he was at Joe's Bar at 11:00 PM on the eleventh

(ii) We observe a picture of the inside of Joe's Bar taken at 11:00 PM on the eleventh that contains a face that closely matches the face in a picture labeled 'John Doe.'

Call the situation that involves the id-action in (i) o_1 and the situation that involves the id-action in (ii) o_2. Note that both involve a resource situation: in o_1, it is the action of sending the email, and in o_2 it is the action of taking and labeling John's picture. The id-situation in (i) is of type, say, $JohnDow_{OL}$, a type in the online environment, and we assert the identity proposition

(1) $o_1 \models JohnDow_{OL}$

The id-situation in (ii) is of type, say, $JohnDow_{Phy}$, a type in the physiological environment, and we assert the identity proposition

(2) $o_2 \models JohnDow_{Phy}$

Apparently, there is a common referenced event, *ref*, of John Doe being at Joe's Bar at 11:00 PM on the eleventh. If we judge that *ref* is indeed a common referenced event, we thereby fuse $JohnDoe_{OL}$ and $JohnDoe_{Phy}$, forming the super-type, say, *JohnDoe*, and o_1 and o_2 are then taken as evidence for *ref* by way of propositions (1) and (2).

Now suppose the eleventh was a Saturday, and suppose we have the following.

(i′) We observe an email from `John.Doe@Acme.com` that states that he was at Joe's Bar at 11:00 PM on the fourth

(ii′) We observe a picture of the inside of Joe's Bar taken at 11:00 PM on the eighteenth that contains a face that closely matches the face in a picture labeled 'John Doe.'

Call the situation that involves the id-action in (i′) o_{1a} and the situation that involves the id-action in (ii) o_{2a}. Both involve resource situations similar to those in o_1 and o_2, respectively. The referenced event for o_{1a}, call it ref_1, is John being in the bar at 11:00 PM on the fourth, and the referenced event for o_{2a}, call it ref_2, is John being there at 11:00 on the eighteenth. We assert the following.

(1′) $o_{1a} \models JohnDow_{OL}$

(2′) $o_{2a} \models JohnDow_{Phy}$

Id-situations (i′) and (ii′) may be further evidence, by way of propositions (1′) and (2′), for fusing $JohnDow_{OL}$ and $JohnDow_{Phy}$. Fusing the identity types produces an additional correspondence between the referenced event *ref*, ref_1, and ref_2 in that they involve the same identity type. These referenced events (and hence, indirectly, id-situations (i), (ii), (i′), and (ii′)) may provide evidence regarding John Doe's habit of visiting Joe's Bar Saturday evenings. A large number of diverse id-situations relating to a single, fused identity may allow us to build up not just a history of an individual but even a full-blown identity.

As another example, suppose we do the following.

(iii) We determine that the DNA from skeletal remains found beneath a car park in Leicester match the DNA from the blood on a garment known to have been worn by Richard III.

(iv) We read in a chronicle that Richard III was hastily buried at a certain location near medieval Leicester.

Call the id-situation that involves the id-action in (iii) o_3 and the id-situation that involves the id-action in (iv) o_4. Note that o_4, but not o_3, involves a resource situation, namely, the situation in which the chronicler wrote about Richard III's burial. The id-situation in (iii) is of type, say, $RichardIII_{Phy}$, while that in (iv) is of type, say, $RichardIII_{Soc}$, and we assert

(3) $o_3 \models RichardIII_{Phy}$

(4) $o_4 \models RichardIII_{Soc}$

Relevant to what is of interest here, o_3 relates to a referenced event, call it ref_3, to do with the demise of the last monarch of the House of York in a certain time interval and his remains being interred at a given location. Id-situation o_4 relates to a referenced event, call it ref_4, to do with the same monarch being buried near a certain city on a particular day in 1485. If we fuse $RichardIII_{Phy}$ and

RichardIII$_{Soc}$, giving the super-type, say, *RichardIII*, we thereby relate ref_3 and ref_4, in effect inferring a new, more specific event, say, ref_{34}, that (unlike ref_3) involves a specific day and (unlike ref_4) involves a specific location. Where e is an event, let prop(e) be a statement in some sufficiently well-defined language whose sense is e. Then we have prop(ref_{34}) $\Rightarrow$ prop(ref_3) $\wedge$ prop(ref_4), where $\Rightarrow$ and $\wedge$ are the usual truth-functional material implication and conjunction connectives. Also note that the referenced event in (iv) helps us rule out such things as a reburial.

9 Category-theoretic View of Identity Fusion

For a systematic account of the interrelation and fusion of identities, we would like something like Barwise and Seligman's information channels. The "parts" would be the identity environments, and an infomorphism f from classification A to classification B would at least include the type function $f^{\wedge}: typ(A) \longrightarrow typ(B)$. The token function, however, is problematic, as we shall see. The core C of such an "identity channel" would compose the types of the various identity environments (not necessarily as a sum), and it would support the following. Suppose that id-situation o_0 in environment E_0 is of identity type T_0 and that id-situation o_1 in environment E_1 is of identity type T_1. If the translation T_0' of T_0 entails the translation T_1' of T_1 in the theory Th(C), then o_0 being of type T_0 carries the information that o_1 is type T_1.

The problem is that not every token (i.e., id-situation) in one "part" (i.e., environment) can be related to some id-situation in some fixed other environment, nor is there a guarantee that an id-situation in one environment relates to at most one in the other environment For example, perhaps I inspect an email from **`Ed@Acme.com`** to **`Al@Acme.com`**. There is no reason to expect any id-actions relating to Ed's or Al's physical aspects; on the other hand, there might be several such id-actions. Resource situations and referenced events do not help here since they might simply be missing from the record for a given environment or duplicated with trivial changes.

We are forced to accept not functions between sets of tokens but rather relations between sets of tokens: for a given identity environment (i.e., classification) A, where C denotes the core, the relation $f^{\vee}: tok(C) \leftrightarrow tok(A)$ in general relates some elements in *tok*(C) to multiple elements in *tok*(A) and relates some to no element in *tok*(A) (when the id-situations are "missing"). This requires us to relax the sum construction. The "core" is no longer unique, that is, the construction is no longer universal. We have, then, a *cocone*, but not a universal cocone (which is a sum). The token relations have to be constructed ad hoc for each environment. Some indication of how this is done was given in the above discussion, where, for various cases, we referred to various combinations of the id-situations, resource situations, and referenced events. The type function can remain a total function $f^{\wedge}: typ(A) \rightarrow typ(C)$ and must be such that the type of an id-situation in a given environment is reflected in its type in the core. We call such a function-relation pair $(f^{\wedge}, f^{\vee})$ an *id-map*, denoted by undecorated $f: A \rightarrow C$. Instead of an information channel, we might refer to an identity channel.

Each identity proposition c |= γ in the core C can be given a confidence indicating that some identity proposition in one identity environment is evidence for some identity proposition in another environment. Suppose we have environments A and B with id-maps f: A → C and g: B → C. Suppose also that c |= γ in the core has a certain confidence $p \in [0, 1]$. Let $a = g^{\vee}(c)$, $b = f^{\vee}(c)$, and $g^{\wedge}(\alpha) = f^{\wedge}(\beta) = \gamma$. Then a |= α carries the information that b |= β with confidence p (and vice versa). (Fig. 3 shows a variation with domains A_i, A_j, and A_k, associated identity propositions a_i |= α_i, a_j |= α_j, and a_k |= α_k, respectively, and id-maps f_i, f_j, and f_k, respectively.) Confidence in a core proposition is less than certainty since tokens corresponding to the connection token are "missing" in some environments. Note, however, that some environments may contribute more tokens, which would generally increase confidence.

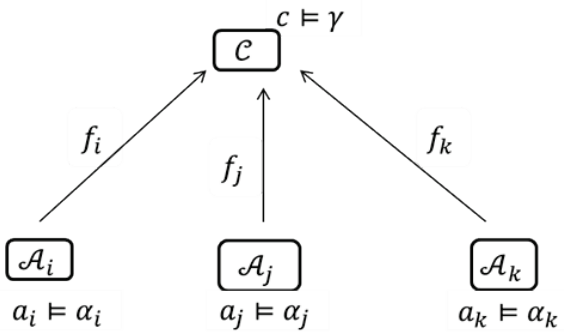

Fig. 3. An identity channel with corresponding identity propositions for the "parts" and the "whole" (core). The confidence in the core proposition is the confidence that a proposition in a part carries the information expressed by any of the corresponding propositions in the other parts.

As mentioned, the partition of identity activity into environments is somewhat arbitrary. In particular, there is a general-to-specific hierarchy of environments. Fusion of identity types, therefore, is often a multistage process as one moves up this hierarchy on multiple branches. Indeed, in our sketches of three environments, we noted fusions within the environments. Since what are parts in Fig. 3 could themselves be cores in subsumed identity channels, we would typically associate a confidence less than 1.0 with the identity propositions associated with these parts.

Fusing types generally reduces confidence in identity propositions since fusion is an inductive process. On the other hand, fusion generally increases confidence in the identification of the players in the referenced events and asserted identities. For example, fusing the identity types relating to the email fom and picture of John Doe in Section 8 results in an identity proposition with confidence less than 1.0 but increases the confidence that we have put our finger on a particular person. The theory associated with the core of an identity channel can be thought of as merged constraints that pick out a noumenal individual (a sort of genotype) behind the id-situations that arise (collectively forming various phenotypes) from its interactions with the various environments. As we fuse more and more of the identity types, these constraints progressively narrow the field. Such increase in one aspect accompanied by a decrease in a mirror aspect is characteristic of Galois connections in category theory.

10. Conclusion and Future Work

We presented preliminary work on a framework for constructing identities that takes as its point of departure Barwise's channel theory, founded on category theory, and allied work on situation semantics and situation theory. The framework takes actions relating to identity, both judgments and assertions, as fundamental. Actions are events (with an agent), and events are basically situations, although situations can include larger regions of space and time and generally accommodate context. Identity-related situations are classified under (identity) types, which can be fused. Identity-related activity is partitioned according to the 'form of life' involved into (identity) environments, similar to Barwise's classifications. An identity-related action is often based on another identity-related action as a resource, and all these actions may reference events in an individual's life. The result is a layered picture of identity-related situations, which is easily accommodated in situation theory. The category-theoretical constructs of Barwise's channel theory, however, must be relaxed when identity in multiple environments is considered. We explained the reasons for this and sketched what needs to be done to recover the category-theoretic foundation.

Identities in this framework retain references to the situations from which they are constructed, so provenance, context, and narrative detail are retained. The role of the "spectator," as well as those of the "subject" and "sample" (data or evidence), are acknowledged (Stevenage et al. 2012). The point of view here is that of an investigator assembling and analyzing evidence for identities. At this point, however, we do not focus on how identities may be discovered or constructed but rather focus on the structure of the evidence and the structure of the process by which identities may be induced from the evidence. Category theory is an ideal tool here since it is ultimate mathematics for representing and reasoning about structure.

We note that how individuals construct their identities in online settings has become an active area of research (Poletti and Rak 2014). We also suggest that our framework, being neutral regarding how identities are discovered or constructed, can be useful across a range of interests, including personal and group identity as addressed in psychology and the social sciences (Vignoles et al. 2011).

References

Awodey, S., 2010. *Category Theory* (2nd Ed.). Oxford University, Press: Oxford, UK.

Barr, M., Wells., C., 1990. *Category Theory for Computing Science*. Prentice Hall: Upper Saddle River, NJ.

Barwise, J., and Perry, J., 1984. *Situations and Attitudes*. Cambridge, MA: MIT Press.

Barwise, J., Seligman, J., 1997. *Information Flow: The Logic of Distributed Systems*. Cambridge Tracts in Theoretical Computer Science (44), Cambridge, UK.

Devlin, K., 1991. *Logic and Information*. Cambridge University Press: Cambridge, UK.

Fiadeiro, J.L., 2005. *Categories for Software Engineering*, Springer: Berlin.

Flavell, J. H., 2004. Theory-of-Mind Development: Retrospect and Prospect. *Merrill-Palmer Quarterly*, 50(3): 274─290.

Goldman, A. I., 2008. *Simulating Minds*: The Philosophy, Psychology, and Neuroscience of Mindreading (Philosophy of Mind), Oxford University Press: Oxford.

Kratzer, A., 2011. Situations in Natural Language Semantics. *The Stanford Encyclopedia of Philosophy* (Fall 2011 Edition), Edward N. Zalta (Ed.), CSLI, Stanford University, Stanford, CA.

Lawvere, F.W., Schanuel, S.H., 2009. *Conceptual Mathematics. A First Introduction to Categories* (2nd Ed.). Cambridge University Press: Cambridge.

MacLane, S., 1998. *Categories for the Working Mathematician*. (2nd Ed.). Springer-Verlag, New York.

Pierce, B.C., 1991. *Basic Category Theory for Computer Scientists*. MIT Press: Cambridge, MA.

Poletti, A., and Rak, J., eds. 2014. *identity technologies. Constructing the Self Online*. U. of Wisconsin Pr: Madisson, WI.

Simmons, H., 2011. *An Introduction to Category Theory*. Cambridge University Press: Cambridge, UK.

Stevenage, S. et al., 2012. Annual Report 2012, Super-Identity Project, University of Southampton. Southampton, UK, http://www.southampton.ac.uk/superidentity/

Vignoles, V. L., Schwartz, S. J., and Luyckx, K., 2011. Introduction: Toward an Integrative View of Identity. *Handbook of Identity Theory and Research.* Schwartz et al. eds. Springer.

Wittgenstein, L., 1991. *Philosophical Investigations: The German Text*, with a Revised English Translation (by G. E. M. Anscombe). Wiley-Blackwell 3rd ed: Hoboken, NJ.