Assembling the ECIpedia: Refining concepts in context

David McDonald

SIFT LLC., 114 Waltham St., Suite 24 Lexington, MA 02421 USA

Mark Burstein

SIFT LLC., 114 Waltham St., Suite 24 Lexington, MA 02421 USA

James Pustejovsky

Brandeis University, 415 South St., Waltham, MA 02453 USA

JAMESP@CS.BRANDEIS.EDU

DMCDONALD@SIFT.NET

BURSTEIN@SIFT.NET

Abstract

We are compiling a repository of minimal, broadly applicable concepts for the benefit of performers in DARPA's Communicating with Computers (CwC) program. Dubbed 'Elementary Composable Ideas' or ECIs, these concepts are broadly applicable cross-domain concepts that underly many lexical semantic terms. As the basis for the ECI repository, we are incorporating and adapting ontological and lexical resources that appear in various forms in TRIPS, VerbNet, Hobbs' commonsense axioms and others sources, and using ideas from DOLCE, Generative Lexicon, and object-oriented knowledge representation to create a common representational framework that supports composition and interpretation in different contexts using the same ECIs. After motivating this effort and looking at the sources we are incorporating, we illustrate the kind of abstraction and conceptual composition we hope to achieve by looking at selected examples in the treatment of space and time as regions and the specialization of preposition interpretation through composition.

1. Introduction

We are developing a repository of concepts and their inferential consequences – in collaboration with and to be shared among the groups participating in DARPA's "Communicating with Computers" (CwC) program (DARPA, 2015), where the focus is on bi-directional communication in multi-modal dialog and determining the meaning of communications relies heavily on the context provided by the environment shared by the participants and the history of their communications with each other. A key element of the research program is the notion that there is a set of perhaps a few thousand basic ideas tied to words like "*up*" that can be used in a wide variety of contexts and understood in each context to denote things tied to this basic systems of ideas, presumably learned during childhood, even as they are composed with others to take on specific meanings..

This paper describes the work done thus far toward defining a representational framework for these notions, called *Elementary Composable Ideas* or 'ECIs', and building an on-line repository for their collection.

1.1 Communication in context

Communication relies on language is designed to incorporate context (Piantadisi et al. 2012), so much so that we should view the 'message' conveyed by an utterance as a combination of its word-ing plus the context in which it is made.

The need for context is the most obvious in short utterances. Matthew Stone (2005) illustrates this with the utterance "*I would like coffee*." Depending on the specific context in which it is used, the word "coffee" might be referring to the drink, a flavor of ice-cream, or a sports team with that name. In CwC, the favorite example is "*add one more*." It also takes its meaning from the ongoing situation. For example, imagine there is a child or a computer program building things out of blocks. Figure 1 depicts two situations where we are in the middle of different building projects. If the context is the arch on the left, then "add one more" means to build another arch. Alternatively, if the context is the two stacks on the right, it means to make another stack. If the context were both the arch and the stacks, then "add one more" would probably mean to add the lintel block over the two stacks. Since utterances like this will generally only occur within a dialog, any uncertainty about what the speaker means can be cleared up by asking what interpretation was intended. But the basic ECI notion of "add entities to collections/structures/regions" is conveyed as the starting point for each contextually grounded interpretation.



Figure 1. Context provided by either the scene on the left or the alternative on the right

The question is how it is that such a simple phrase can take on a range of particular meanings depending on the context in which it is uttered. In this example, the information carried by the literal content (ignoring context) is only that there is something that can be extended, 'added to', and that whatever that is, it already has some things 'in' it since we are being told to add 'one more' to it. What the extendable set or region or structure is and what we are to add is strictly dependent on what is salient in the situation we are in. It could be another piece of fish to be placed on our fishmonger's scale. It could be another cup of flour into a bowl when baking. It could be one more thing placed on our Amazon wish list so we can have enough in our cart to qualify for free shipping. It could be add another measure to the score we are composing, create another character for the game we are playing, increasing the number of people on our dinner reservation etc. Moreover, a person would never say anything this 'elliptical' unless they were certain that the person they are addressing shares their context and is fully aware of what they intend.

We believe that it is the norm rather than the exception for an underspecified term to only get a meaning when its is used in a context. This is closely related to what Situation Semantics refers to as *the efficiency of language* (Barwise & Perry, 1983) – the fact that exactly the same phrase can

be used over and over, with the same literal meaning, but with differing interpretations when used in differing situations. This is a defining aspect of language, whereby we are able to deploy a fixed, relatively small stock of words to express an unbounded set of complex ideas and intentions, or, as Humboldt said "to make infinite use of finite means" (1836). The question, of course, is how is this done.

1.2 Design principles

The design principles we are using in developing the representation language for the ECI repository we are building are intended as a step toward answering this question. They also motivate our decision early on to not just take some existing system off the shelf and just make minimal modifications. As we will illustrate and elaborate in greater detail in the course of this paper, we are following these principles.

- The concepts included are minimal and mostly 'elementary' in the sense that they are not readily defined by logical composition, and are not specific to a domain.
- The taxonomic lattice makes extensive use of multiple inheritance as one means of composition, allowing complex concepts to be created through the composition of simpler ones to the extent that if concepts share roles, they also share the relations, processes and other inferences involving those roles.
- Concepts have associated descriptions of what follows or should be assumed when they apply: their implicatures. When the concept applies in different generalized contexts, the concept may alter the implicatures as they apply differently to those contexts.
- These constitutive descriptions are inherited and refined as concepts are specialized and can be contributed by any of the concepts that were compose to create the current one.
- All such contraints are context specific. The set of states in which a particular constraint or implicature holds is an essential part of their definition.

2. Representation

In this section we lay out the 'ingredients' that have gone into our repository so far, and introduce our formalized representation of an 'elementary composable idea' (\$2.2) and how the concepts we have represented this way are organized. We then introduce our notation for concepts and the rationales behind it (\$2.4).

2.1 Combining existing repositories

No large concept repository that we are aware of fits all of out criteria for facilitating refinement in context while providing a basis of inference (see §5). In particular we believe that any account of refinement in context needs to incorporate ideas from Generative Lexicon Theory (GL) especially co-composition (Pustejovsky, 1995, 2012). Nevertheless we of course did not start from scratch. Instead we drew from the work of long-time collaborators who are also performers in the CwC program. Figure 2 lists the present contributors; the thickness of the connecting lines reflects how much each pair has been addressed thus far.



Figure 2. The groups and repositories contributing to the ECIpedia

This is a summary of what has gone into our ECIpedia repository at this point.

- We worked out a mapping of the thematic roles used in VerbNet and in the TRIPS logical form to a consensus set based largely on the choices made by the ISO Semantic Roles group (ISO 24617-4: see Bunt & Palmer 2013).
- We developed an alignment between the classes and inheritance structure of the TRIPS ontology and the semantically consistent classes of VerbNet, focusing on basic notions that appear as verbs or events, identifying their counterparts in these resources, and then augmenting them with content primarily about their defining process characteristics as needed for GL.
- We developed an initial set of ECIs for the "Blocks World" domain that is one of the conversational domains of the program, using GL/VoxML representations developed at Brandeis as examplars. (Pustejovsky & Krishnaswamy, 2016).
- We worked with USC/ISI to integrate a set of predicates developed by Jonathan Gordon and Jerry Hobbs for their commonsense reasoning axioms into the repository, and began using them to specify the relationships found in various generalized processes as parts of event ECIs.

2.2 Elementary Composable Ideas

One of the premises of the CwC program is that we reason with compositions of abstract conceptual units that take on a meaning when there is a active context that people can access to get the corresponding particulars. On this view, language, even when it is augmented by gesture and facial expression. is a low-bandwidth channel, and what we pass through this channel are instructions for how to compose complex ideas from simpler ones in a given context.

To refer to these simplest concepts Paul Cohen coined the term "*elementary composable idea*," or ECI. In part, he chose this neologism as a deliberate strategy to avoid established notions about how much information is packed into individual concepts or what form the represented information might take, but it was also chosen because the research focus was on naturally occurring minimalist communication that could only be properly interpreted in context. The term implies that there is something fundamental and basic, 'elementary,' about ECIs, perhaps because they were learned during our earliest experiences as children interacting with our world, and certainly because of the human tendency to communicate enough information to allow the communicative intent to be recognized given the shared context, but less than would be required to compose the full meaning strictly from the words used. The other key aspect of ECIs that we see in the name is that they are designed to be composed. The complex concepts of everyday life (blocks, chairs, professors) are taken to be composed from elementary ones, though, again, the literal meanings of these compositions is only enough to identify shared situations in prior and present experience.

In the remainder of this paper we will use the term 'ECI' when we are talking about the properties and use of the representations forming elements of our repository, and use 'concept' when the technical detail is not important.

2.3 Organizing the ECI repository

As will soon become clear, inheritance plays a large role in the determination of active aspects of the concepts in our repository. This has the effect of placing a premium on the design of the uppermost concepts, the 'upper model.' We have constructed our upper-model by adopting the on-tological stance taken by DOLCE (§3.1), and using selected abstract concepts from our contributing repositories, notably from Jerry Hobbs (Gordon & Hobbs, 2017).

- The upper model defines a relatively small set of abstract categories such as eventuality, entity, attribute, configuration, region, time, space, scale, individual, collection, parts, shape, dimension, sequence, change, opposite.
- The next level consists of the simple (elementary) ECIs. Elementary categorizations of concrete things we could encounter in the world such as basic physical objects; event classes such as transition/movement; attributes and relations like color, weight, height, speed, location, orientation, force, containment, etc.
- Composite ECIs (CCIs) are compositions of ECIs and other CCIs. The composition is frequently done as part of reading a text (§3.6). Alternatively, composite categories that we want to reify in the repository are created by composition through multiple-inheritance in their definition. Simplifying enormously, we could for example define 'running' as the composition of movement and high-speed (§3.4.1).

When the distinction is not relevant we will refer to all of these simply as ECIs.

2.4 Formalism

Our design choices for the formalism for representing the contents of ECIs and their interrelationships in our ECIPedia repository were influenced by the KL-One tradition (Woods & Schmolze, 1992) and the description logics descended from KL-One such as OWL (McGuinness et al., 2004), as well as object-oriented programming languages such as the Common Lisp Object System (CLOS) (Gabriel et al., 1991). We treat ECIs as intensional concepts in a T-box (T = terminological). Instances of ECIs (e.g., to reflect a particular interpretation of a state of affairs) are individuals that populate an A-box (A = assertional).

We have a Lisp-inspired notation for our representation which is convenient for defining and revising ECIs. When we update the content of the ECIpedia website, we run a conversion program to produce the equivalent JSON expressions that the web server uses to layout the page content. Our own 'live' instantiation of ECIs for our dialog system is in Lisp. Others in the CwC program have developed implementations in Python.

ECIs are organized into a single inheritance hierarchy. Except for the ECI at the top of the lattice, every ECIs inherits from one or more other ECIs. In the example shown just below, the ECI named 'transfer-location' inherits from a ECI named 'transfer'.

```
(def-eci transfer-location (transfer)
  :comment "the activity of something moving between locations"
  :properties (:trips ont::move :vn (motion-51* put-9*))
  :args ((@theme :isa entity)
                      (@initial-location :isa location)
                           (@final-location :isa location))
        :const
  ((:holds-in @start-e (location-of @theme @initial-location))
        (:holds-in @e (move :theme @theme))
        (:holds-in @end-e (location-of @theme @final-location))))
```

Every ECI has a set of *fields* of different sorts There is typically a :comment field containing a string that describes what the ECI is intended to mean or how it would be used for the benefit of a human reader (and fellow developers). A :properties field records what the ECI corresponds to in other repositories. In this example, it notes that transfer-location maps to the 'move' concept in the TRIPS ontology, and to two particular classes in VerbNet.

The arguments field (:args) in this example defines three arguments that identify the participants in this eventuality, and restricts their allowed types. Arguments are inherited. ECIs lower in the hierarchy can narrow the restrictions stated on ECIs they inherit from. The constitutive field (:const) lists the propositions that hold or the actions that happen in each of three standard *states* inherited from the ECI for event: before (@start-e), during (@e) and after (@end-e), following Pustejovsky's theory of event structure (1991). The propositions are stated in terms of arguments, either those defined locally on the ECI or inherited from more general ECIs. Note that in the 'move' statement, its thematic role, which is indicated by :theme, is bound to this ECI's theme argument, indicated as @theme.¹

3. The value of abstraction and composition

The purpose of our repository is to explore the semantic basis of the efficiency of language. For example, how do we formalize our intuition that the 'up' of "*pick up the block*" spoken in the blocks world means the same thing when we say "*move the notes up a fifth*" when talking about a musical score.

We are not trying to assemble another encyclopedic fact base or almanac to support question answering or competitions like the Winograd Challenge (Levesque et al., 2011). Our effort has much in common with the OntoNotes project where people recognized that the sense distinctions made in WordNet were too fine-grained for practical use (Palmer et al., 2007) and developed an alternative by grouping senses together (Hovy et al., 2006). We are trying to design a small set of basic concepts (ECIs), that can be deployed in different configurations and composed to provide effective semantic interpretations in task-based dialogs where a rich situational context is always available.

In this section we begin by looking at the major influences that underlie our approach to semantic analysis. We begin by looking at the basis of our upper model (\$3.1) and introducing additional ECI fields that are crucial for modeling context (\$3.2). We then consider the relationship of ECIs to words (\$3.3.1) and the abstraction and simplicity that comes from positioning region high in the taxonomic lattice (\$3.5). We close this section looking at prepositional relations including their functional interpretations (\$3.6).

3.1 Upper Model

For our upper model – the concepts at the very top of the taxonomic lattice – we follow the ontological perspective developed for DOLCE (Gangemi et al., 2002). It takes a *descriptive* (rather than revisionary) ontological stance, which takes the structure of language and commonsense to have ontological relevance. This permits us to model thing that people can conceive of yet don't exist in the actual world (unicorns, fairy rings in the garden, the present king of France).² This also entails using a 3D approach where events are situated in time which is an independent dimension. It is *multiplicative* rather than reductionist. For example it permits two different things to occupy the same space if they are of different sorts, such as a statue and the clay that was used to make it. DOLCE comes with a full axiomatization of its concepts and ontological positions in terms of 1st order formulas and as Common Logic (Masolo et al., 2003).

In addition to adopting DOLCE's commonsense view of events, we also adopt their treatment of attributes and their values. Simplistic views of attribute value as terms can be subject to paradoxes,

^{1.} We prepend an atsign (@) character on every argument to get around a limitation in JSON. Thematic role arguments, in particular, will often have the same names as useful ECIs ('instrument') which would lead to a clash in JSON if the symbols that represent them were not distinct.

^{2.} A revisionary ontology such as BFO will adjust the set of concepts to fit what is presently accepted as a scientifically accurate structure of reality and considers linguistic and cognitive issues as secondary.

such as Barbara Partee's observation that saying "the temperature is 90" and "the temperature is rising" and assuming the validity of substitution of equals, lets you conclude that "90 is rising", which is of course nonsense. DOLCE instead distinguishes the value of an attribute on a particular individual, say a rose, from what it corresponds to in the quality region for colors. To say that the rose has faded is to say that its color qualia now points to a different point in the color region.

3.2 Context: Habitats and Affordances

Our representation for objects employs the account of logical polysemy of Generative Lexicon theory where every lexical concept embodies multiple, equally available functional aspects, its *qualia structure*. The qualia enable different affordances. A coffee mug, for example, via it's formal qualia, affords the ability to act as a paper weight by warrant of being a physical object of a certain weight and size. Its functional affordance (its telic) is to hold coffee for you to drink.

We formalize other, more contingent and context sensitive affordances and implicatures as 'habitats', illustrated just below (Pustejovsky, 2013; McDonald & Pustejovsky, 2014). This ECI for a sheet of paper is modeled as an artifact that has three different sets of affordances depending on what habitat it is in. Two that are defined locally in the ECI and one (h-3d) that is inherited from physical-object. In any situation you can treat a piece of paper as an object which affords holding and moving it. The paper is readable in most alignments provided it is facing you. It affords writing on it when it is flat on a surface.

```
(def-eci paper-sheet (physical-artifact)
 :habitats
   ((h-paper-readable :intrinsic
      (align @self-y @env-y)
       (surface-direction @self-front -z-vector))
     (h-paper-writable :intrinsic
       (align @self-top @env-y)
       (surface-direction @self-front -z-vector)))
:affordances
  ((:when-in h-3d
       (grasp :theme @self)
       (hold :theme @self)
       (move :theme @self))
   (:when-in h-paper-readable
       (read :theme (symbols-on :ground @self)))
    (:when-in h-paper-writable
      (write :destination
               (surface-of :ground @self)))))
```

3.3 Thin concepts

We conjecture that our repository will have the greatest utility and will be usable as-is to the largest set of people, if its concepts, elementary and composite, are 'thin' – each one contributing just a small amount of information.

To explore this, we can look at what goes into representing the concept that the word *up* refers to. This concept, which we can call 'up', is a 'direction'. Obvious directions are points of the compass like 'north' or 'southeast', but more generally, directions determine the paths that things

follow when they change (change of location while walking, change in amount when making a stack of blocks).

The natural grounding of 'up' is as the opposite of the direction established by the force of gravity (i.e. 'down'), but other sorts of things provide the requisite gradient where 'up' is applicable: musical notes on a scale, temperatures, wind – in short, any sort of quantity whose possible values lie on a gradient of some sort can be characterized as being 'up'.

The concept of 'up' is subservient on the concepts of 'direction', which in turn is subservient on a particular aspect of the phenomena of change. If we continued, this regression would get progressively more abstract and of less utility in a shared ontology because fewer 'operational level' concepts would be able to benefit from its associated implicatures. Instead we have to stop somewhere and decided that some concepts should be taken as primitive. In this instance we could call our stopping point – our simple ECI – *directed change*.

3.3.1 The status of words

An ontology is distinct from a lexicon. The names of the concepts in an ontology could in principle be replaced with arbitrary generated symbols (G0037), except that it would make unmanageable demands on the memory of the people who maintain it, which is why we use ordinary words as concept names. But names are not words, no matter how hard that is to keep straight given the concept names we tend to use.

There can be a *realization* relationship between concepts and the lemmas of the lexicon. In our own implementation, ECIs are mapped to lexicalized TAG trees so the lemma includes the relevant syntactic patterns. Indeed, any ontology that is used as the target representation in natural language understanding will be 'lexicalized' in some fashion.

The natural tendency by ontologists to use ordinary words can be problematic. For example, if we look at particular cases that we characterize as employing the concept 'up', we quickly find that we want to call them different things. We talk about, e.g. water *rising*, or the National Debt *increasing*, and as soon as we consider comparisons there is *higher* or simply *more*.

The question is whether these alternatives constitute different concepts or whether they are simply synonyms. Of course whether true synonyms actually exist is a matter of considerable debate (see, e.g. Miller et al. (1990)). But from an semantic perspective we should ask whether there are substantive difference in presuppositions or implicatures between these alternatives, or whether the differences in usage are actually determined by the semantic field, differences in connotation simply matters of habit and best accounted for using a distributional analysis.

3.4 Composition in definitions and interpretations

3.4.1 Defining concepts through composition

As we just defined it, the minimal conceptualization of up is just that it is a 'direction' whose opposite is 'down'. Given 'up' we can now define the concept climb by just having it inherit from both 'move' and 'up'. 0

(def-eci climb (move up))

The concept 'move' includes a the path along which to travel, and 'up' supplies the needed value.

There are many variations on climb that we call *ascend* or *rise* or even *increase* or *expand*. We take it as axiomatic that co-composition is a ubiquitous aspect of language interpretation that makes it unnecessary to distinguish many of these notions at the conceptual level. The meaning of a verb can not be determined without simultaneously taking into account what arguments it has, what's moving and what are they going 'up' on: ballons expand, bread rises/expands, you climb/ascend a ladder or cliff face, you move up a note on a score, stock prices increase/rise/ascend. If there are significant semantic differences in, say, the presuppositions of these different formulations, then the concept of climb should be specialized to accommodate them. If that it not the case then the type restriction of the arguments can be noted as part of the interpretation rules for how the concept is realized given a different arguments, and the result will be the same basic move+up ECI.

3.5 Regions as abstract composites

Taking our cue from Dolce where the concept of region is very broadly defined, we have positioned our composite ECI for region very high in our taxonomy. Here is its full current definition.

```
(def-eci region (abstraction composite-entity)
  :comment "The concept of 'region' is taken abstractly here using
  the notion of demarcation or boundary over a domain to identify
  or 'locate' particular parts of it.
  You can have a region of the number line, a region of a
  process (e.g. S-phase in cell division), a region of space.
  Regions have (possibly fuzzy) borders and interiors in which other
  things can be located, as well as supporting the notions of relative
  location, extent (size), distance, neighborhood, and such."
    :args ((@interior :isa region)
        (@components :isa collection) ;; of regions, 'members'
        (@extent :isa measurement) ;; 'size', 'length', ...
        (@dimension :isa dimension) ;; 1D, 2D, ...
        (@containing-region :isa region)) ;; for borders
    :const ((exists @interior)))
```

At this level all we are saying about a region (in its constitutive) is that it has an interior. By modifying the assertions in the const field and adding adding or specializing arguments we can, e.g., define a point as a region with no containing region (i.e. it is atomic) and no extent. A one-dimensional region is the composition of region and the ECI for one-dimension. We define two and three dimensional regions similarly with each adding an additional measurement argument.

Notice that at the level of region that nothing is said about what kind of things its components are or what kinds of stuff its extent is measuring. Indeed, we define both space and time as specializations of region. For example here is how we define a interval of time.

Notice how the borders of an interval are now time points. This permits us to define the Allen relations in a natural way. Given an ECI temporal-order over figure and ground arguments that are constrained to be intervals, we have simple definitions in terms of the relationships of the interval's borders (i.e. their beginning and ending time points).

```
(def-eci allen-before (temporal-order)
    :const ((earlier @figure.end @ground.begin)))
```

This does require defining primitive temporal predicates (ECIs) for same-time, earlier, and later. But from there it is a short step to defining the spatial relations of the Region Connection Calculus 8 (RCC-8, Randell et al. (1992)) in terms of the 2D borders of spatial regions by adding a predicates for 'touch' and 'overlap'. The predicates for 'inside' and 'disconnected' are already defined at the level of the region ECI.

3.6 Prepositional relations

While the mechanics will vary from parser to parser, the semantic interpretation of a free preposition (not bound to the verb like "pick up") is a two-place relation governed by the preposition. We follow Talmy (1975, 2000) in taking these prepositional relations to be relationships between a figure and a ground. They are relations between a figure that is an entity (or an eventuality) and grounds that are regions, though note that at this level we are not committing to a particular sort of region (spatial, temporal, color, nationality, ...). By staying abstract and leaving it to later context to determine the choice we hope to increase the applicability of the ECI.

For basic topological prepositions like "at" (Herskovits, 1986), the preposition itself carries very little information. The essential meaning of "at" is that the ground is the location of the figure. In terms of ECIs that comes to this:

This works equally well for *the meeting at MIT* and *the meeting at noon*, the first establishing a location in space and the second a location in time. This is a simple illustration of how the character of the relation between figure and ground (is it in space or in time) is determined by the properties of the ground,

3.6.1 "Add one more" in terms of collections and regions

Recall our examples at the beginning of the paper about how the underspecified literal meaning of "*add one more*" can take on any number of full interpretations depending on what the context is. Adding something to an existing collection involves changing its location (if a is persistent physical

object) or possibly creating it for the target location. Before we add it – say we're adding a log to a fire – the log is not a member of the collection (of logs in the fire), afterwards it is. But is the log *in* the fire? (Compare that to adding a log to the pile of logs next to the fire.)

If we take a fire to occupy a region of space (at some level of granuarity) then something, such as a log, can be placed in that region. In this sense, the log will be in the fire when the fire can be construed as a region or container, in the same way that "in the water" describes partial or whole containment within the region occupied by the body of water. A pile of logs usually has fuzzy borders, so the logs that are already part of the pile can provide a support for another log, which would licence the preposition *on*, as well as *in*.

3.6.2 Functional spaces and the influence of context

Further examples illustrate that the preposition *at* is more than a mereotopological relation, but can invoke the functional space for an activity. For example, being *at the piano* or *at the blackboard* presuppose a particular orientation and configuration of the figure relative to the ground. There are constraints on the figure (the person in this case) that are imposed by the ground (piano, blackboard). They follow from what it means for the figure to exploit the affordances provided by the ground: playing the piano in the first case, writing on the board in the second. In both cases, the person must assume a particular orientation. They must facing towards the functional side of the object and be close enough to it to operate it. Moreover, the figure (person) should be attending to the object, focusing on it. A person leaning on the back of the piano might be focused on it but they cannot play it, though they would be in a good position to accompany the pianist as a singer.

These are facts about pianos and blackboards, which we could capture in their ECIs or, if we find that these facts are common to whole class of objects, could be expressed as part of a super-ECI that they inherit from. The question is what can be done during the interpretation of the prepositional relationship that will 'convey' this information about the constraints to the figure from the ground ECI on which it is stated. The solution depends on the GL machinery described earlier that we use to develop the ECIpedia and on the use of multiple inheritance. The shared facts about orientation when in a figure/ground functional locative relation can be encoded in this ECI.

The ECIs 'imposes-configuration' and 'attend-to' are mixins rather than 'base' ECIs with linguistic realizations. They add (or merge) their information about possible habitats and what they afford to the information on the base ECI. In this case the mixins are designed to interact with any semantic interpretation that involves the figure/ground relation, but particularly for prepositions like 'at' or 'on' that have locative force.

The semantic interpreter of the prepositional relation imposes a context that identifies which constituent is the figure and which the ground. When it applies the constraint expressions of the ground, this triggers the 'functions-at-ground' state in the habitat of imposes-configuration. That in turn causes the statements in the corresponding affordance to be asserted, adding the implicature that the person, in our example, can be assumed to be facing the piano.

It is also the case that a person 'at' a piano can usually be assumed to be sitting down on something of the correct height so that their fingers can touch the keys. It is important to decide how to judge how much detail about default assumptions of this sort should be included. So far the criteria we use is just an intuition about how frequently something occurs and how standard it is. Here the idea that an activity assumes a canonical orientation of the performer's intrinsic long axis applies also to sitting, standing, running, swimming, and sleeping.

A similar issue is how to decide how much to lump together in the habitats or affordances of a given multi-use mixin ECI. That would appear to come down to whether the candidate for possible 'lumping' ever applies independently. For the concept of 'attending' to something (whose ECI is essentially identical to imposes-configuration modulo the implicature it asserts), it should be its own mixin ECI because it does apply in different circumstances than just those that invoke a canonical orientation. Attending or concentrating on something applies to studying 'for' an exam, listening attentively 'to' music or 'to' a person you are talking with.

4. The ECIpedia project

Early on in the CwC program, an *ECI committee* was formed by researchers in who had resources to contribute. They were charged with clarifying and elaborating the notion of an elementary composable idea.

- Martha Palmer at UCBoulder: VerbNet (Kipper et al., 2000; Palmer et al., 2017)
- James Allen at Rochester and IHMC: The TRIPS ontology (Allen et al., 2008, 2013)
- Jerry Hobbs at USC: Formalization of Common Sense (Hobbs & Gordon, 2005; Gordon & Hobbs, 2017)
- James Pustejovsky at Brandeis: Generative Lexicon, VoxML (Pustejovsky & Krishnaswamy, 2016)
- Mark Burstein at SIFT: leads the effort to assemble the ECIpedia repository, and designed the initial representation system.

4.1 Repository and Web site

Doing the actual merging fell on our shoulders at SIFT, along with the mandate to manage a program-wide repository of ECIs. Figure 3 is a screen from the web site we have stood up. Presently the web site, https://ecipedia.sift.net, requires credentials to view.³ Once it is reasonably mature, our plans are to publish the repository that backs the ECIpedia on GitHub and use its push request machinery for negociating updates.

^{3.} Username: cwc password: communicate

D. MCDONALD, M. BURSTEIN, AND J. PUSTEJOVSKY



Figure 3. Screenshot from the ECIpedia showing the detail page for 'put'.

4.2 Consumers

One of the central goals of the CwC program has been to provide the AI community with a legacy that will enable future programs and individual efforts to be able to build on the shoulders of the knowledge representation accomplishments of CwC. The ECIpedia is the first step on a path towards this, and one of its most important roles is to permit CwC performers to share their results and compare them to other groups' treatments.

This is a two-way street. For example, we have been working with people in the 'Musica' project (Quick & Morrison, 2017). They are working to enable dialogs (jam sessions) and joint composition between a computer and human musician. Part of this involves telling the machine how you want to modify a score. They looked at an early release of the ECIpedia and found gaps that we have collaboratively fixed. The phrase "*the first three notes*",' for example, led us to refine our treatment of specifying selection from sets. As an early consumer they also pointed out some flaws in the ergonomics of the ECIpedia website.

We are also starting to work with the animated story-telling group at SRI Princeton and the neural modeling group at the University of Washington. We fully expect our ECIpedia to change in

content and usage patterns as it 'comes in contact' with a wider and wider set of actual or potential users and contributors.

5. Similar work

There are of course other large concept repositories than just the ones we drew on for the ECIpedia. These are basically of three different sorts.

- Omnibus systems are a combination of massive, linked-data "fact repositories" with entries in the tens of millions and comprehensive, lexicalized knowledge bases that are created by mining Wikipedia, WordNet and many similar machine-readable resources: Freebase (Bollacker et al., 2008), Yago (Hoffart et al., 2013), DBpedia (Lehmann et al., 2015), BabelNet (Navigli & Ponzetto, 2012).
- Fact-finders are systems that mine the whole web for facts, such as the Never-ending Learning project (Mitchell et al., 2015), or TEXTRUNNER and its successors (Etzioni et al., 2011).
- Knowledge bases and Ontologies We include here 'traditional', usually curated systems such as Cyc (Lenat et al., 1990), ConceptNet (Liu & Singh, 2004), or particularly FrameNet (Baker et al., 1998). Other possibilities are GUM (Bateman et al., 1995), and SUMO (Pease et al., 2002)).

These repositories are invaluable in their target applications, such as question answering, but as we described earlier (§3), the ECIpedia is not intended to be as comprehensive as these systems. They are not designed to provide the stock of context dependent minimal intensional concepts that we need for the ECIpedia. In some ways FrameNet is most similar in that there is an implied emphasis on associations between these basic words and ideas and the frames or situations that help define their usage in different contexts, and give rise to the set of implications that are warranted in the specific context at hand. Specifically, we have focused on these basic characteristics in our modeling effort:

- COMPOSITION we are emphasizing the identification of simple ideas that have realizations in different contexts, and whose interpretations can be made specific only when they are composed with others ideas or used in a discourse grounded in a specific context. Think of "up" and then of some expression that use it and consider how much context and how much of the additional implication for the phrase comes from those grounding expressions.
- RELATIONSHIP OF SEMANTICS TO PROCESSES Many implications or expectations associated with these basic concepts come from their relationships with processes and change. We seek to model the processes associated with both events and entities (e.g., their affordances) as a means of making explicit how they relate to roles in processes and their associated temporally relative states.
- HABITATS AND LINKED AFFORDANCES that facilitate making concepts' implicatures sensitive to different contexts.
- MIXIN CONCEPTS that let us bundle useful behavior even though the mixin has no linguistic realization and can lie outside the main line of the type lattice.

6. Concluding remarks

It was clear that, though we sought to emphasize different things, the only sensible approach to arriving at a comprehensive, shared set of ECIs was to try to build on, merge and adapt as many useful pre-existing resources as we could. To that end subgroups within this group looked at mappings between the various representations they were using to try to identify shared concepts, and appropriate generalized inferences. Among them,

- James Pustejovsky (Generative Lexicon) has been working with Martha Palmer to align GL representations with those of VerbNet, which has resulted in some changes to VerbNet to improve their alignment (Brown et al., 2018), and make them more consistent with our ECI formalism.
- James Pustejovsky's team at Brandeis also worked with Bruce Draper and Ross Beveridge at Colorado State to define ECI representations for gestures in GL/VoxML format for their gesture-oriented dialog system.
- James Allen's team at IHMC worked on aligning their TRIPS ontology with VerbNet, and clarifying the meaning of their choice of thematic roles (Allen & Teng, 2018).
- Mark Burstein and David McDonald worked with TRIPS and GL-based ECI representations to develop ECIpedia, demonstrate interpretation in specific dialog contexts, and to incorporate relations and inferences from Hobbs' commonsense reasoning formalization into the ECIpedia framework.

This is very much a work in progress. These and other groups are contributing the concepts that they are need to interpret interactions in the dialog systems they are developing. We will continue to expand both the catalog of ECIs and refine the representation system to support additional forms of contextual composition as we go forward. Ultimately the mark of success will be if the formalism and ECI library continues to be useful after the CwC program is finished, and in the extent to which these ideas and content are used by other people engaged in cognitive modeling.

Acknowledgements

This work was supported by Contract W911NF-15-C-0238 with the US Defense Advanced Research Projects Agency (DARPA) and the Army Research Office (ARO). Approved for Public Release, Distribution Unlimited. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

References

Allen, J., de Beaumont, W., Galescu, L., Orfan, J., Swift, M., & Teng, C. M. (2013). Automatically deriving event ontologies for a commonsense knowledge base. *Proceedings of the Tenth International Conference on Computational Semantics* (pp. 23–34).

- Allen, J., & Teng, C. M. (2018). Putting semantics into semantic roles. *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics* (pp. 235–244).
- Allen, J. F., Swift, M., & de Beaumont, W. (2008). Deep semantic analysis of text. Proceedings of the 2008 Conference on Semantics in Text Processing (pp. 343–354). ACL.
- Baker, C. F., Fillmore, C. J., & Lowe, J. B. (1998). The Berkeley FrameNet project. Proceedings of the Seventeenth International Conference on Computational Linguistics (pp. 86–90). ACL.
- Barwise, J., & Perry, J. (1983). Situations and attitudes. Cambridge, MA: MIT Press.
- Bateman, J. A., Magnini, B., & Fabris, G. (1995). The generalized upper model knowledge base: Organization and use. In N. J. I. Mars (Ed.), *Towards very large knowledge bases: Knowledge building & knowledge sharing*, 60–72. Amsterdam: IOS Press.
- Bollacker, K., Evans, C., Paritosh, P., Sturge, T., & Taylor, J. (2008). Freebase: A collaboratively created graph database for structuring human knowledge. *Proceedings of the 2008 ACM SIG-MOD International Conference on Management of Data* (pp. 1247–1250). ACM.
- Brown, S. W., Pustejovsky, J., & Palmer, M. (2018). Integrating generative lexicon event structures into VerbNet. *Proceedings of the Eleventh International Conference on Language Resources and Evaluation*.
- Bunt, H., & Palmer, M. (2013). Conceptual and representational choices in defining an iso standard for semantic role annotation. *Proceedings of the Ninth Joint ACL-ISO Workshop on Interoperable Semantic Annotation*. Potsdam.
- DARPA (2015). DARPA-BAA-15-18: Communicating with Computers.
- Etzioni, O., Fader, A., Christensen, J., Soderland, S., & Mausam, M. (2011). Open information extraction: The second generation. *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence* (pp. 3–10). Barcelona: IJCAI.
- Gabriel, R. P., White, J. L., & Bobrow, D. G. (1991). CLOS: Integrating object-oriented and functional programming. *Communications of the ACM*, 34, 29–38.
- Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., & Schneider, L. (2002). Sweetening ontologies with DOLCE. *Proceedings of the 2002 International Conference on Knowledge Engineering* and Knowledge Management (pp. 166–181). Springer.
- Gordon, A. S., & Hobbs, J. R. (2017). A formal theory of commonsense psychology: How people think people think. Cambridge, MA: Cambridge University Press.
- Herskovits, A. (1986). Language and spatial cognition: An interdisciplinary study of the prepositions in English. Cambridge, MA: Cambridge University Press.
- Hobbs, J. R., & Gordon, A. S. (2005). Toward a large-scale formal theory of commonsense psychology for metacognition. *Proceedings of the AAAI Spring Symposium: Metacognition in Computation* (pp. 49–54).
- Hoffart, J., Suchanek, F. M., Berberich, K., & Weikum, G. (2013). YAGO2: A spatially and temporally enhanced knowledge base from Wikipedia. *Artificial Intelligence*, *194*, 28–61.
- Hovy, E., Marcus, M., Palmer, M., Ramshaw, L., & Weischedel, R. (2006). OntoNotes: The 90%

D. MCDONALD, M. BURSTEIN, AND J. PUSTEJOVSKY

solution. *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers* (pp. 57–60). Association for Computational Linguistics.

- Humbolt, W. v. (1836). On the diversity of human language construction and its influence on the mental development of the human species. Cambridge, MA: Cambridge University Press.
- Kipper, K., Dang, H. T., Palmer, M., et al. (2000). Class-based construction of a verb lexicon. Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence (pp. 691–696).
- Lehmann, J., et al. (2015). DBpedia A large-scale, multilingual knowledge base extracted from Wikipedia. *Semantic Web*, *6*, 167–195.
- Lenat, D. B., Guha, R. V., Pittman, K., Pratt, D., & Shepherd, M. (1990). Cyc: Toward programs with common sense. *Communications of the ACM*, *33*, 30–49.
- Levesque, H. J., Davis, E., & Morgenstern, L. (2011). The Winograd schema challenge. *Proceedings of the AAAI Spring Symposium on Logical Formalizations of Commonsense Reasoning* (p. 47).
- Liu, H., & Singh, P. (2004). ConceptNet: A practical commonsense reasoning tool-kit. BT Technology Journal, 22, 211–226.
- Masolo, C., Borgo, S., Gangemi, A., Guarino, N., & Oltramari, A. (2003). WonderWeb deliverable D18: Ontology library. Technical report, Laboratory For Applied Ontology, Trento, Italy.
- McDonald, D., & Pustejovsky, J. (2014). On the representation of inferences and their lexicalization. *Advances in Cognitive Systems*. Cognitive Systems Foundation.
- McGuinness, D. L., Van Harmelen, F., et al. (2004). OWL web ontology language overview.
- Miller, G. A., Beckwith, R., Fellbaum, C., Gross, D., & Miller, K. J. (1990). Introduction to WordNet: An on-line lexical database. *International Journal of Lexicography*, *3*, 235–244.
- Mitchell, T., et al. (2015). Never-ending learning. Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence.
- Navigli, R., & Ponzetto, S. P. (2012). BabelNet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *Artificial Intelligence*, *193*, 217–250.
- Palmer, M., Bonial, C., & Hwang, J. D. (2017). VerbNet: Capturing english verb behavior, meaning and usage. In S. Chipman (Ed.), *The Oxford handbook of cognitive science*. Oxford University Press.
- Palmer, M., Dang, H. T., & Fellbaum, C. (2007). Making fine-grained and coarse-grained sense distinctions, both manually and automatically. *Natural Language Engineering*, 13, 137–163.
- Pease, A., Niles, I., & Li, J. (2002). The suggested upper merged ontology: A large ontology for the semantic web and its applications. *Working Notes of the AAAI-2002 Workshop on Ontologies and the Semantic Web* (pp. 7–10).
- Piantadosi, S. T., Tily, H., & Gibson, E. (2012). The communicative function of ambiguity in language. *Cognition*, 122, 280–291.
- Pustejovsky, J. (1991). The syntax of event structure. Cognition, (pp. 47-81).

Pustejovsky, J. (1995). The generative lexicon. Cambridge, MA: MIT Press.

- Pustejovsky, J. (2012). Co-compositionality in grammar. *The Oxford handbook of compositionality*, *371*, 382–397.
- Pustejovsky, J. (2013). Dynamic event structure and habitat theory. *Proceedings of the Sixth International Conference on Generative Approaches to the Lexicon* (pp. 1–10). ACL.
- Pustejovsky, J., & Krishnaswamy, N. (2016). VoxML: A visualization modeling language. *Proceedings of LREC*.
- Quick, D., & Morrison, C. T. (2017). Composition by conversation. *arXiv preprint arXiv:1709.02076*.
- Randell, D. A., Cui, Z., & Cohn, A. G. (1992). A spatial logic based on regions and connection. Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (pp. 165–176). Cambridge, MA.
- Stone, M. (2003). Knowledge representation for language engineering. In A. Farghaly (Ed.), *A handbook for language engineers*, 299–366. Center for the Study of Language and Information.
- Talmy, L. (1975). Figure and ground in complex sentences. *Proceedings of the Annual Meeting of the Berkeley Linguistics Society* (pp. 419–430).
- Talmy, L. (2000). Toward a cognitive semantics, volume 2. Cambridge, MA: MIT press.
- Woods, W. A., & Schmolze, J. G. (1992). The KL-One family. *Computers and Mathematics with Applications*, 23, 133–177.