

## Learning concepts and language for a baby designer

**Madan Mohan Dabbeeru and Amitabha Mukerjee**

*Indian Institute of Technology Kanpur, India.*

*We introduce the “baby designer enterprise” with the objective of learning grounded symbols and rules based on experience, in order to construct the knowledge underlying design systems. In this approach, conceptual categories emerge as abstractions on patterns arising from functional constraints. Eventually, through interaction with language users, these concepts get names, and become true symbols. We demonstrate this approach for symbols related to insertion tasks and tightness of fit. We show how a functional distinction - whether the fit is tight or loose - can be learned in terms of the diameters of the peg and the hole. Further, we observe that the same category distinction can be profiled differently - e.g. as a state (clearance), or as a process (the act of insertion). By having subjects describe their experience in unconstrained speech, and associating words with the known categories for tight and loose, the frequencies of words associated with these can be discriminated. The resulting linguistic labels learned show that for the state profile, the words “tight” and “loose” emerge, and for the action, we get “tight” and “easy”. Once an initial grounded symbol is available, it is argued that knowledge-based systems based on such symbols can be sanctioned by its semantics, as well as its syntax, leading to more flexible usage.*

### Symbols and Design Reasoning

Machine design systems have been used for encoding the final design, and for downstream functions such as analysis or manufacturing. In the attempt to generalize it to conceptual design, it is tempting to define a set of symbols and rules for modeling a domain (knowledge-based systems) [5, 6, 7, 11, 12, 17 and 30]. However, the word “symbol” as used in the context of

computers has a far narrower interpretation from that in human usage, which can lead to considerable inflexibility.

In computers, symbols are defined formally, i.e. only in terms of other symbols, and lack the connection to domain experience underlying flexible usage among human designers. If we may present an analogy, computer usages of symbols are like the understanding of a colour symbol like “red” by a blind man; he knows that it is an instance of something called “colour”, and that “green” and “blue” are other colours, and maybe even that “crimson” and “vermilion” are shades of “red”, but his understanding is dramatically different from that of a sighted person, because the semantics is not connected to direct experience.

When human designers use symbols, their usage is flexible and even for very abstract terms, the semantics is well-grounded and they can easily come up with detailed instances of the idea. Computers using symbols may be able to provide instances for basic symbols (based on programmer definitions), but not for symbol compositions. Further, the semantics defined by programmers cannot take into account many contextual factors that may change the interpretation of a symbol; indeed, rules that apply over a general domain often need to be modified to fit the problem - “general rules never decide concrete cases” [28]. While symbolic reasoning systems have been used successfully in the context of design for very limited situations, we argue that they may prove difficult to scale up for the following reasons:

- Designers often differ widely in what they mean by any term; the meaning of any term is rarely independent of its context. Thus attempts at defining a “standardized” vocabulary may not fructify.
- Formal semantic models, typically based on “intension”, i.e. a set of rules defined on other formal symbols (e.g. [24]) provide a very narrow, inflexible interpretation.
- Cognitively, related terms are often organized in a loose, hierarchy that is defeasible - i.e. memberships can be overruled in exceptional situations (a “bird” may be an animal that can fly, but it may also be a toy resembling the animal). On the other hand hierarchies in computational models (ontologies) are rigid, leading to unforeseen failures in novel situations.
- Crucially, symbols defined in terms of other symbols alone are ungrounded, like the blind man’s “red”. This implies that every poss-

ible relation with other symbols, and all possible consequences for actions must be explicitly encoded. Thus, it must know that after a wing injury, a “bird” may no longer fly, but still remain a bird. The number of such axioms is potentially unbounded. Grounded symbol semantics avoids this problem because the model of birds that don't fly, or the roles of wings in flight would also have emerged at some stage of experience, leading to a graded inference (“birds fly” is a rule, but it can be overruled).

### **Learning Symbols**

Here we propose to treat the term “symbol” in machine design as it is understood by human designers, and not as in formal algebras. Thus a symbol would constitute a close coupling of a term or label and a semantic representation or “image schema” [18]. In the design context, the image schema may be viewed as a set of constraints abstracted from different design experiences. Good designs repeatedly reveal certain inter-relations among the design variables; we propose to learn these as image schemas (e.g. “clearance”, “oscillating motion”, etc.) Such image schemas need to be discovered from functional associations during design exploration, since the functions defining them are often too complex to be modeled by simple rules, and in any event, the definitions may change considerably based on context. This may be why designers find it difficult to define terms they use regularly, and why it is felt that much of their knowledge may be implicit [28]. This discovery process must start with the simplest concepts and build up, which is why we call our approach the “baby designer enterprise”.

In earlier work, we have explored the emergence of image schemas in a baby designer through computational simulations [23]. Here we focus on learning the labels for a schema, so the resulting label-schema pair becomes a true symbol. We demonstrate the process by learning the simple distinction between tight and loose fits. The baby designer is given a set of explicit performance measures and show how graded, grounded schemas can emerge based on this functional distinction.

Subsequently, we consider the baby designer as interacting with a human expert, who describes the different situations using language. In this interaction, we assume no knowledge of language or grammar; all we assume is that the system is that the human narratives is available as word-separated text rather than raw speech data (human infants show good

word-boundary separation skills starting around 9 months of age). Then we show how our computational baby can learn some labels merely by considering the frequency of words that are associated with these conceptual distinctions. Thus, the system now has a symbol in terms of both its label and its grounded semantics. The semantics for it may now broaden with further exposure. This approach has also been attempted in other domains ([14, 27 and 29]).

We observe now that the semantics of a symbol is much more than just its referent. The semantics also encodes many subjective aspects, such as how the referent is being viewed - how it is being *profiled* [19]. For example, entering a space may be profiled either as an action (process, “enter”) or as a whole unit (atemporal relation, “into”). Though the conceptual structure is largely the same, there is a subtle difference caused by the focus on different aspects of the same structure. Thus the process view would accept temporal modifiers, the state view would not. This led us to conduct two different experiments for collecting language narratives. In the first, we ask the users to describe, in unconstrained English, their experience with an already assembled peg and hole. In the second we ask them to describe the experience of putting the peg into the hole. The word associations are significantly different, but one word does appear in both contexts - thus, its semantics is already enriched by these two meanings.




Fig.1 *Profiling*. The same conceptual structure may be viewed as a relational complex or a process (action). Langacker [9].

At the end of this process, the semantics of these words is grounded in one limited context, so to our baby, the words mean have this limited meaning, but as its experience broadens, she will generalize the meaning to new situations. Note that the semantics learned is not a “primary” or “core” meaning – indeed, whether such privileged senses exist is itself uncertain. The sense learned is just one in a continuum of possible interpretations, many of which will be learned with further exposure, including possible metaphorical extensions to other domains.

### *Are designs emergent?*

The approach here focuses on a situation where the image schema is available before the label is known. For human designers, this assumption may be questioned, since designers (e.g. students) learn many concepts by being told – i.e. the concept arises after its name is given. For a human learning language, beyond a small initial inventory of symbols, the vast majority of words are learned through its correlation with other words [4]. Nonetheless, the early inventory of symbols is crucial, for it provides grounding for the compositions that define later symbols. Only in this manner would grounding be available for the new concept. In the design scenario, the need for experiencing a domain directly is all the more crucial, which is the basis for hands-on approaches in design pedagogy, as opposed to other didactic disciplines.

A second reason for symbols to emerge in a design context is because the designer faces a challenge far greater than mere problem-solving or search, since the very constituents of his problem are ill-defined. One of the first tasks the designer must do is to discover representations at the right level for encoding and formulating the problem. Some of these representations, if they keep recurring, may become symbols. This process is related to the discovery of *chunks*, an abstraction formed from many input variables, and a topic often associated with design expertise. For example, we note that a designer who is an expert in a particular domain is “confident of immediately choosing a good [design] based on experience” [13]. Clearly, some sort of simplification of the problem space has occurred for the expert. While models of expertise have focused on chunks as they arise among trained designers, the process applies to all learning, and forms the backbone of symbol learning from our “baby designer” to the very best designers. This is why we feel this mode of learning is scalable -- from the very early stages demonstrated in this paper, to far more complex situations encoding large swathes of domain knowledge.

### **Related work: Discovering patterns in design spaces**

An early attempt at discovering patterns in the design space of shapes may be seen in relation to 2D shapes [25]. Another approach to discovering chunks in design operates within the tradition of formal ontologies [22]. Here a learning layer, operating as a manager (M-agents), is added to a system being used to create designs. The M-agents consider the good

designs that have come up, and try to identify some patterns which eventually become chunks that are added to memory. Further, the effectiveness of a new chunk can be tracked in subsequent designs to ascertain its utility. In design problem re-formulation mechanism suggested by Sarkar et al [26], designers can identify latent relationships among different design variable groups by using Singular Value Decomposition on the co-occurrent variable matrix. This helps designers to redefine the design problem by re-representing it with a possible reduction in dimensionality.

None of these proposals however learn the semantics underlying the symbol in a grounded manner, using the feedback received from exploring the space of “good designs”. An optimization-based approach towards this problem can be seen in [8]. Here the results of multi-objective optimization are analyzed manually to propose the possible inter-relationships between design variables necessary for “good design”. The approach proposed here automates this process for both linear and non-linear relationships. Also, the discovered “dimensions” are proposed as “chunks” that are the putative basis for symbol discovery.

### **Baby Designer Enterprise**

Our approach differs from the above in that it discovers structures in the input space, and posits these as proto-symbols, for which labels are then determined through language association. Our “baby designer” is at the apprentice level that has no prior knowledge about domains, though it knows many machine learning algorithms and has a bias towards shorter, information-condensing representations. It is given a set of raw descriptors (design variables) and a set of performance metrics defined on these variables so it can be evaluate different design instances. By exploring the design space it first identifies the pattern of good solutions, the Functionally Feasible Regions or *FFRs*. Next, it seeks to determine if the *FFR* is embedded in a lower-dimension sub-space, reflecting some constraints that holds among design variables among the “good” solutions. These inter-relations are proposed as one of the mechanisms for discovering chunks in design.

Discovering chunks results in a set of priors which encode domain-specific knowledge, but these are not the same as symbolic rules, because they are not explicit. For example, the system cannot provide a justification for such decisions. We may compare these with decisions that a human designer knows are good but finds difficult to justify, e.g. by saying

“looks right” [1]). However, if similar chunks are observed repeatedly, especially in different domains, one may become conscious of the pattern, a process called *reification*. These reified chunks are more stable, and are sometimes called perceptual symbols in cognitive science [2]; in this work we refer to these as image schemas.

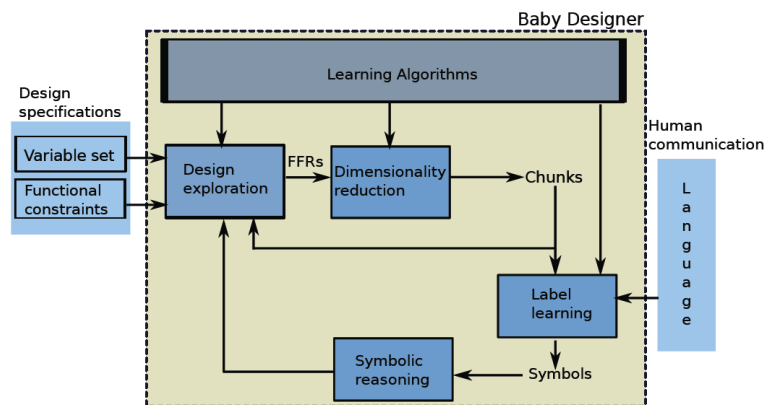


Fig. 2. *Architecture of the Baby Designer*: The baby designer learns patterns as in an apprenticeship situation. It is presented a design problem with design variables and a set of functional constraints that map design instances to performance metrics. While exploring the design space, it uses an inventory of domain-general learning algorithms to discover patterns that hold among the better designs, which may become chunks if they recur often enough. Labels for these may be learned after exposure to language, when they become true symbols. Implicit associations (priors) become encoded as domain rules in this symbolic space, thus enabling symbolic reasoning.

Subsequently, if the system interacts with a language user and finds that a string is strongly correlated with the image schema, then it may identify this string as the label, and the resulting structure may become a true symbol, with both a semantics (the image schema) and a linguistic label. Any priors that were earlier implicitly known will now get mapped explicitly into symbolic rules. These symbolic structures are situated, in the sense that they map the semantics in the “right” context at the “right” level of detail. Eventually, this will enable the system to reason with the symbol more flexibly, generate exemplar instances, and to justify its decisions in symbolic terms.

### Learning Containment

Now we present the computational baby designer with the task of

inserting a peg into a hole. A very early discovery for human infants is that pegs must be narrower than holes, i.e. the hole-width  $w$  must be greater than the peg thickness  $t$  (Fig. 3a). For the human designer, such primitives lie at the heart of the semantics for many symbols relating to containment, assembly, dimensioning, fit, etc. Thus, learning design symbols cannot start as an adult - it must start with our earliest conceptual achievements, which is why we call this the “baby designer enterprise”.

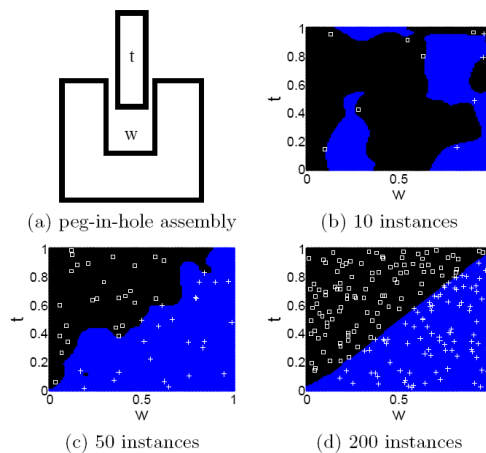


Fig. 3. Learning through experience that inserted-object-must-be-smaller-than-container ( $w > t$ ). After a few instances, the experience is unstable, but the pattern converges after sufficient instances.

The key to this learning is that functional criteria must be available. In these initial stages, we consider the baby designer as an apprentice, so that functional criteria are given by a mentor or some other external source. In this first learning task, the functional criterion is that the peg must enter the hole. Those instances where it can enter (+) constitute the FFR, which is distinguished from failures ( $\square$ ). The system has a function generalization algorithm that can generalize the pattern in the parameter space- here we use a back-propagation perceptron.


We observe from Fig. 3 how, after experiencing just a few instances (Fig. 3b), the pattern is inchoate, so the baby keeps trying to insert some fat pegs into smaller holes, but this exploration itself keeps filling up the negative (black) area of the figure (Fig. 3c,d). Eventually the defining boundary becomes sharper, and at some point it does not change as much



with new experiences, so that it feels it may have discovered a stable pattern, at least implicitly. This knowledge can be thought of as a simple prior defined on the  $w, t$  space, an implicit version of the rule that  $w$  must be greater than  $t$ .

### Learning about clearance

In the next step, let us consider that our baby designer has is exploring different types of successful insertions - particularly, those involving instances of tight fit and loose fit. Given a functional description of these, it learns the corresponding FFRs (Fig. 4a, b, FFRs in gray). These functions also become sharper with greater experience.



**Fig.4.** Emergence of chunks for fit: “tight” vs. “loose”: Given a task which requires a “tight” fit, the Functionally Feasible regions (FFRs) learned after 20 instances is diffuse (a), but improves within 100 trials (b). For loose fits, a well-learned stage is shown (c). These FFRs are modeled using PCA, resulting in a 1-D characterization where the principal eigenvector represent invariance in  $w-t$ . Thus, when considering the concept of clearance, the number of parameters involved reduces from  $(w; t, 2-D)$  to the emergent chunk  $w-t$ , a single parameter (1-D). This learning process results in two chunks,  $C_T$  and  $C_L$  but the system does not have any names for these yet.

Next, it attempts to see if the good instances may be lying along some

low-dimensional sub-space of the design space  $\langle w, t \rangle$ . To begin with, we may use linear dimensionality reduction. Trying out Principal component analysis [3], we find that the results on the early learning after 20 samples are not very clear, but after 100 samples, the first eigenvalue is clearly dominant (33.60, 0.11), and the first eigenvector (-0.72, 0.69) is along a 45 degree line in the  $w, t$  space (Fig. 4a, bottom). Thus, though the design space had two dimensions, we discover that the distribution of good instances of tight fit is largely distributed along one dimension. A parallel eigenvector is found to be dominant in the loose fit case. These two 1-D lines constitute the basis for the categories tight ( $C_T$ ) and loose ( $C_L$ ). The invariant along either line is the quantity  $w-t$ , which becomes the learned “chunk”; its value eventually may form part of the semantics for the symbol “clearance”.

Such correlations, which are embedded as lower-dimensional manifolds in the high-dimensional design space, may be rather common in design. For example, if strength is to be maximized while minimizing weights, then many dimensions need to be balanced – they would rise (or fall) in tandem. Thus, for good designs, these inter-relations may result in a single chunk. Discovering these interdependences is a first step towards the process of creating semantically rich models of design. While the example here deals with only linear subspaces, we have elsewhere dealt with non-linear manifold discovery [23].


## Language Mapping

At this stage, we have an implicit notion of the categories  $C_T$  and  $C_L$ , but we cannot relate this idea to other concepts because there is no handle or label with which to refer to it. The label is a crucial part of the symbol - without it, the chunk or image schema that has been learned cannot be related to a broad host of other concepts. Implicit rules where these chunks play a role cannot be used explicitly to justify decisions, or to reason about the effects of decisions. More importantly, a label such as “clearance” stabilizes the semantics through social convention; without it the semantics may drift with new, similar experiences.

In order to learn a label, we obtain human commentary on the same  $C_T - C_L$  distinction which has been already learned. We provide human subjects with a simple apparatus - several flat pieces of wood with a hole, and some cylindrical pegs that fit in these holes with different types of clearance (Fig. 5). We give different combinations of these to human subjects and

have them describe their experience with them in unconstrained English. Then we would associate individual words appearing in these descriptions with the concepts, and see if any good labels would emerge.

As discussed above, the same insertion task may be viewed under different profiling distinctions. In some situations the complete task is viewed as a whole (atemporal), while elsewhere one may consider its evolution over time (temporal). We also explore how these distinctions may lead to differences in the symbol used to refer to the same concept.



**Fig.5.** Peg-in-hole assembly : A, B, C with three hole sizes (22.5, 17.1 and 12.74) average diameter respectively) and pegs 1 through 6 (22.4, 21.2, 16.9, 14.5, 12.5, and 10.3) are used. A:1, B:3 and C:5 are tight fits, A:2, B:4 and C:6 are loose.

For the human designer also, profiling distinctions are important. For example, “easy to insert” takes a process view of function, while “loose fit” relates more to an atemporal view, although both may be talking about the same design. In different design contexts, a designer may use one symbol or another and no amount of standardization can do away with such differences. Indeed, this reveals an important aspect of symbol use contrasting with the formal view - a symbol is not merely the object being referred to (referent), or even a class definable by a set of attributes and associations, for it also encodes a subjective view of the referent. As a cognitive linguist has colourfully said, the same person may be referred to as “eminent linguist” and “blonde bombshell” [10], expressions that highlight different profilings for the same referent.

To test the effect of profiling in greater detail, we designed two different scenarios for collecting user data. Although these are not experiments in the traditional sense (there is no hypothesis to be validated), we refer to these as experiments for want of a better term. In the first experiment, we placed already inserted peg-hole assemblies in front of our subjects, and

they were instructed not to pick up the block with the hole, so they could not practice insertion actions. In the second experiment, we gave them the block with the hole and the peg separately and had them play around while inserting them. In either case, we asked them to describe the interaction between the peg and hole in plain, unconstrained English, and did not correct any grammar errors etc, though we did transcribe them into written texts. No constraints were imposed on the language, and at least in one instance a subject held forth at length about the colour and smoothness of the apparatus and other aspects – this was the only narrative excluded from our analysis, though it does not substantially affect the result.

Also, before collecting data, users were permitted two trial rounds, one with a tight fit and one with a loose fit, without telling them anything about these.

### Associating linguistic labels

Next, we outline the process used to discover a linguistic label for the image schema that has been learned. The association of a word with a concept  $C$  can be measured in many ways; the machine translation literature poses many association measures. One of the simplest is based on conditional probability, but here there is some debate re: the direction of causality - should we consider the conditional of word  $w$  given  $C$  or  $C$  given  $w$ ? Let  $n_T$  and  $n_L$  be the number of words in the  $C_T$  and  $C_L$  narratives, and let word  $w$  have a count of  $k_L$  and  $k_T$  in each narrative. Then one may estimate the conditional probability  $p(w/C_T)$  as  $k_T/n_T$ . For the conditional probability of  $C$  given  $w$  we may adopt Bayes' rule, which gives us

$$p(C/w) = \frac{p(w/\tilde{C})p(C)}{p(w)}$$

Now, since the prior word frequency  $p(w)$  is independent, we have

$$\frac{p(C_L/w)}{p(C_T/w)} = \frac{p(w/C_L) p(C_L)}{p(w/C_T) p(C_T)}$$

Furthermore, the number of  $C_L$  and  $C_T$  instances in the training data is roughly the same so  $p(C_L)$  and  $p(C_T)$  are equal. Hence the ratio of the two conditionals are equal in both directions. Thus, if we find the  $w, C$  pair that maximizes this ratio, then either conditional would be maximized.

So our objective is to compute  $\max_i \frac{p(w_i/C_L)}{p(w_i/C_T)}$  so that  $w_i$  would have the strongest association with  $C_L$ . Similarly, the inverse ratio is to be maximized for the strongest association with  $C_T$ .

For this task, the user narratives are first transcribed and all the narra-

tives relating to each fit situation (tight or loose) are combined. The word counts  $n_T$ ,  $n_L$ ,  $k_L$ ,  $k_T$  are used to compute the conditional ratio and the top five correlations in all four Concept-Profile are presented in Table 3. We also observe that subjects use many morphological variations - e.g. for “tight” (count=26), we also have “tighter” (3) “tightly” (4) etc. We may use stemming [15] (discard common affixes) to count only the roots of such words. We contrast below the results with and without stemming – and find that even without stemming, correlations are quite strong.

Another step frequently adopted in NLP is to remove the frequent words which occur in many diverse contexts, so their relevance in a particular task may be less. These include particles and grammatical markers like *the*, *a*, *an*, *of*, *in*, *to*, *is*, *am*, etc. However, we found good results even without removing these words. Correlations discovered without these steps implies minimal assumption of linguistic knowledge for the word association process.

**Experiment: STATE:** The purpose of this first experiment is to focus on state distinctions - i.e. collect the spoken English data for the situation where the subjects is given a peg already inserted into the hole. We then collect their unconstrained English descriptions.

**Method: Apparatus:** six wooden pegs (1...6) and three blocks A,B,C as shown in Fig. 5.

**Participants:** Eighteen IIT Kanpur graduate students, both male and female, of age 18-24, participated in generating narratives. Students had back-grounds in physics, mechanical engineering, biology, electrical engineering, chemical engineering and design. Level of competence in spoken English varied somewhat across the group sentence structures were retained as spoken, even if they were ungrammatical.

**Procedure** Each participant is presented with the following instruction:

*“This is a peg and this is a hole. The peg is already inserted into the hole. Play with the assembly, but please do not lift the block with the hole from the table. Describe the interaction between the peg and hole in English. ”*

Note that no reference was made to the tight-loose distinction; many participants reported on many other aspects such as the relative sizes of the peg and hole pairs, their shapes, the kind of construction, etc. In state-profiled trials a peg pre-assembled into blocks A, B, and C was placed on the table and the subject could not to lift these blocks from the table – thus

they could experience the assembly more as a whole, rather than the insertion task as a process. Each subject was given alternating tight and loose assemblies in the order (A:1) - (A:2) - (B:3)- (B:4) - (C:5) - (C:6). Sample narratives for two (A: 1) and (A: 2) of a speaker is given in Table 1.

**Analysis and Discussion:** The narratives for tight {(A: 1), (B: 3), (C: 5)} and loose {(A: 2): (B: 4) and (C: 6)} cases result in a small sample corpus (1099 words for  $C_T$ , and 904 words for  $C_L$ ). Many words appearing in one

- |    |  |
|----|--|
| A1 | This one is very sticky kind of thing [...] it is not moving in the either of directions and its firmly hold it to the basement and [...] i am not able to rotate or pull it to the any side of block. |
| A2 | this piece of art the later piece can be moved easily in all directions [...] and can be removed out of the block and can be rotated either directions.  |

Table1: Transcribed narrative: State profiled. [...] indicates pause.

set do not appear in the other, but fortunately, the top twenty-five words appearing in either set were also present in the other, so the ratios could be computed for these frequent words.

To ensure the relevance of these words, we tried to ensure that the probability of the word appearing in this context is higher than its prior in general usage; i.e.  $p(w/C)$  should be greater than  $p(w)$ . Priors were estimated from a spoken English corpus (based on TV scripts) with 29 million words [16]. For all the top 25 words the conditional was higher than the prior. Table 2 shows the top five words for the tight corpus ( $C_T$ ), in descending order of the ratio  $p(w/C_T)$  by  $p(w/C_L)$ . Thus, the strongest correlations for  $C_T$  are “tight”, “to”, “into”, “not”, “rotate”. Now, the prior probability of words such as “to” and “not” are several orders of magnitude higher - i.e. these are most likely used in a wide variety of situations. Thus the more appropriate words for this situation are “tight”, “into”, and “rotate”. In the unstemmed case, we find “cannot” and “am” in addition to “tight”, “into”, “to”. The term “tight” is the most likely key word in both stemmed and unstemmed cases. The word “rotate” does not appear without stemming since morphological variants such as *rotating* (5 times) and *rotation, rotations* (1 each) are lost. The word *rotate* itself appears only in the state profile, probably because subjects were instructed not to remove the peg so the only action they could try was to rotate it in the hole

State profiled: [tight] corpus								
Term	$f_T$	$\hat{p}(\frac{w}{C_T})$	$f_L$	$\hat{p}(\frac{w}{C_L})$	$f_{T,L}$	$\hat{p}(w)$	$p(w)(T.V)$	$\frac{\hat{p}(\frac{w}{C_T})}{\hat{p}(\frac{w}{C_L})}$
Without stemming								
tight	26	0.02366	3	0.00331	29	0.01447	0.00004	7.14
cannot	7	0.00637	1	0.00110	8	0.00399	0.000121	5.76
to	33	0.03003	8	0.00884	41	0.02046	0.02810	3.40
into	12	0.01092	3	0.00331	15	0.00749	0.00093	3.29
am	9	0.00819	3	0.00331	12	0.00599	0.001294	2.47
With stemming								
tight	34	0.03094	3	0.00331	37	0.01846	0.00004	9.33
to	33	0.03003	8	0.00884	41	0.02046	0.02810	3.4
into	12	0.01092	3	0.00331	15	0.00749	0.00093	3.29
not	19	0.01729	7	0.00773	26	0.01297	0.00660	2.24
rotate	23	0.02093	12	0.01326	35	0.01747	0.00000	1.58
State profiled: [loose] corpus								
Term	$f_L$	$\hat{p}(\frac{w}{C_L})$	$f_T$	$\hat{p}(\frac{w}{C_T})$	$f_{T,L}$	$\hat{p}(w)$	$p(w)$	$\frac{\hat{p}(\frac{w}{C_L})}{\hat{p}(\frac{w}{C_T})}$
Without stemming								
loose	27	0.02983	4	0.00364	31	0.01547	0.00003	8.2
much	10	0.01105	2	0.00182	12	0.00599	0.00117	6.07
can	28	0.03094	12	0.01092	40	0.01996	0.00912	2.83
quite	8	0.00884	4	0.00364	12	0.00599	0.00018	2.43
easily	10	0.01105	7	0.00637	17	0.00848	0.00002	1.73
With stemming								
loose	30	0.03315	4	0.00364	34	0.01697	0.00003	9.11
much	10	0.01105	2	0.00182	12	0.00599	0.00117	6.07
can	28	0.03094	13	0.01183	41	0.02046	0.00912	2.62
easy	17	0.01878	9	0.00819	26	0.01297	0.00023	2.29
in	15	0.01657	13	0.01183	28	0.01397	0.00966	1.4

Table 2. State profiled: [tight] and [loose] corpora.

**Experiment: ACTION:** In this experiment the subjects are permitted to actively insert the peg into the hole, using the same apparatus as above.

A different set of eighteen graduate students, both male and female, of age 18-24, studying physics and mechanical engineering, participated in generating narratives. Each participant is presented with the following instructions:

*“This is a peg and this is a hole. The peg can be inserted into the hole. Describe the interaction between the peg and hole in English. ”*

Peg-in-hole Assy.	Description in spoken English
B3	this is another hole and this is another peg [...] this is also going very tightly [...] it is not going completely inside the hole.
B4	this is another one this is very loose [...] in this hole and we can pass it through the hole very easily

Table 3. Transcribed narratives : Action profile.

As in the previous experiment, each subject is given alternating tight and loose fit situations, and asked to describe their experience in spoken English (sample narratives in Table 3).

Action profiled: [tight] corpus							
Term	$f_T$	$\hat{p}(\frac{w}{C_T})$	$f_L$	$\hat{p}(\frac{w}{C_L})$	$f_{T,L}$	$\hat{p}(w)$	$\frac{\hat{p}(\frac{w}{C_T})}{\hat{p}(\frac{w}{C_L})}$
Without stemming							
tight	26	0.02050	4	0.00462	30	0.01406	4.43
first	13	0.01025	3	0.00347	16	0.00750	2.96
not	30	0.02366	8	0.00925	38	0.01782	2.56
a	11	0.00868	3	0.00347	14	0.00656	2.5
i	30	0.02366	9	0.01040	39	0.01828	2.27
With stemming							
tight	33	0.02603	4	0.00462	37	0.01735	5.63
first	13	0.01025	3	0.00347	16	0.00750	2.96
not	30	0.02366	8	0.00925	38	0.01782	2.56
i	30	0.02366	9	0.01040	39	0.01828	2.27
of	27	0.02129	11	0.01272	38	0.01782	1.67
Action profiled: [loose] corpus							
Term	$f_L$	$\hat{p}(\frac{w}{C_L})$	$f_T$	$\hat{p}(\frac{w}{C_T})$	$f_{T,L}$	$\hat{p}(w)$	$\frac{\hat{p}(\frac{w}{C_L})}{\hat{p}(\frac{w}{C_T})}$
Without stemming							
easy	13	0.01503	2	0.00158	15	0.00703	9.53
second	10	0.01156	2	0.00158	12	0.00563	7.33
easily	18	0.02081	4	0.00315	22	0.01031	6.6
can	13	0.01503	4	0.00315	17	0.00797	4.76
very	21	0.02428	11	0.00868	32	0.01500	2.8
With stemming							
ease	35	0.04046	6	0.00694	41	0.01922	5.83
second	10	0.01156	2	0.00231	12	0.00563	5.00
can	13	0.01503	6	0.00694	19	0.00891	2.17
very	21	0.02428	11	0.01272	32	0.01500	1.91
diameter	13	0.01503	9	0.01040	22	0.01031	1.44

Table 4. Action Profile: [tight] and [loose] corpora: Top five words by conditional ratio.

The top five word associations for  $C_T$  are presented in Table 4 sorted by the ratio of conditional probabilities. It is observed that both with and without stemming, “tight”, and “first” are most relevant associations for action profiled  $C_T$ , though after stemming, the term “tight” emerges stronger. For loose fits when profiled as a process, “easy” is found to have the strongest association both with and without stemming.



**Discussion:** These experiments demonstrate that finding the linguistic labels for image schemas like tight and loose is possible without much prior knowledge, of either the domain or of language. Initially we had thought we may need to use standard techniques like stemming and stop word removal, and the results - that “tight” and “loose” came up so readily in the state profiled experiment - were a pleasant surprise. Further, in the action profiled case, the word for describing the action where the clearance is high, in English, is not “loose”, but “easy” - an association that some may find reasonable. The other interesting fact is that the same word, “tight” is used as the label for two related but somewhat different semantics - the tight fit case profiled as a state, and also an action. This is not very surprising, since languages must encode, in a finite inventory of units, the unbounded phenomena present in the universe.

So this type of polysemy, which is sometimes called lexical polysemy (as distinct from accidental polysemy or homonymy), is extremely widespread, e.g. the word “motor”, which may indicate the engine of a car or an electric motor. Without taking contextual cues into consideration, symbol meanings are impossible to define.

## Conclusion

This work aims to introduce a novel mechanism for symbol acquisition, but the present demonstration is necessarily very limited, as indicated by the “baby designer” appellation. A baby's symbols are initially overspecialized, and apply only in the limited contexts where they learn it. Similarly, the baby designer's symbols will initially have narrow, specific meanings. However, once the initial symbols are available, one may merge other experiences sharing the same label, to broaden the semantics. How such experiences would be amalgamated into a more general conceptual structure remains a key question for further work.

At this stage, our goal is to present grounded acquisition of symbols as an alternative to traditional models used in knowledge-based systems for design. While we have argued that such a system is likely to be more scalable and more flexible in its deployment, clearly the range of actual problems that can be handled by such an approach will remain rather limited for some time to come. The demonstration here is extremely simple, yet the ease and lack of prior domain assumption argues for scalability.

As of now, the system learns patterns only in the apprenticeship mode, which works only for well-understood domains. The symbols learned thus can be used to create flexible knowledge structures for these known domains. Knowledge with a rich semantic basis is a pre-requisite for more difficult challenges such as novel domains or creative designs – but here again, much work remains to be done.

More than the results itself, we feel that the key contribution of this work is to open up several new avenues for further integration of cognitively motivated approaches into computational systems for design. One of these is that of broadening the semantics of an initial symbol via new exposures, a problem we are exploring now. Another question that arises is the compositions of symbols. New symbols often arise as compositions – “thick flange”, “airy hallway”, “pinwheel arrangement” etc. instantiate or blend some aspects from each of its constituents, and which to choose and which to ignore remains a serious challenge for any model of semantic composition. Learning image schemas in novel situations beyond the apprenticeship scenario is another serious challenge.

Another potential contribution of this work is a possible approachment between formal or systematic approaches to design [5, 12] and other, more fluid, emergent views [26, 25]. The symbols that are learned here are initially emergent, but may eventually be used to formalize design knowledge and lead to coherent theories for different domains. This may provide a basis for unifying the so-called systematic and creative camps in design theory.

When will the baby designer be able to solve actual design problems? Clearly, the system will need to be exposed to a lot more situations (equivalent to a fifteen-year childhood, and then a five-year apprenticeship, say), but we will also have to develop much of the learning structures needed to consolidate these experiences. The paradigm proposed here makes only a small start, and like the interaction of a baby with its world, it opens up more questions than answers. We hope that with further work, we may discover the contours of some of these answers, and that these may then illuminate the potentialities of this approach over the coming years.

## References

1. Ahmed, S., Wallace, K. M., and Blessing, L. T. Understanding the differences between how novice and experienced designers approach design tasks. *Research in Engineering Design* 14, 1 (February 2003), 1–11.
2. Barsalou, L. Perceptual symbol systems. *Behavioral and Brain Sciences* 22 (1999), 577–660.
3. Bishop, C. *Pattern recognition and machine learning*. Springer, 2006.
4. Bloom, P. *How children learn the meanings of words*. MIT Press Cambridge, MA, 2000.
5. Bohm, M. R., Stone, R. B., and Szykman, S. Enhancing virtual product representations for advanced design repository systems. *Journal of Computing and Information Science in Engineering* 5, 4 (2005), 360–372.
6. Campbell, M. I., Cagan, J., and Kotovsky, K. Agent-based synthesis of electro-mechanical design configurations. *Journal of Mechanical Design* 122, 1 (2000), 61–69.
7. Chakrabarti, A., Sarkar, P., Leelavathamma, B., and Nataraju, B. S. A functional representation for aiding biomimetic and artificial inspiration of new ideas. *AI EDAM* 19, 02 (2005), 113–132.
8. Deb, K., and Srinivasan, A. *Innovization: innovative design principles through optimization*. Tech. Rep. Kanganal:2005007, Indian Institute of Technology Kanpur, Kanganal, IIT Kanpur, 2007.
9. Ericsson, K. *Expertise*. MIT Press, 1999.
10. Evans, V., and Green, M. *Cognitive linguistics: An introduction*. Edinburgh University Press, 2006.
11. Gero, J. S., and Fujii, H. A computational framework for concept formation for a situated design agent. *Knowledge-Based Systems* 13, 6 (2000), 361 – 368.
12. Gorti, S. R., and Sriram, R. D. From symbol to form: a framework for conceptual design. *Computer-Aided Design* 28, 11 (1996), 853 – 870.
13. Gross, M. D. *Design as Exploring Constraints*. PhD thesis, Department of Architecture, Massachusetts Institute of Technology, February 1986.
14. Guha, P., and Mukerjee, A. *Language Label Learning for Visual Concepts Discovered from Video Sequences*. Springer-Verlag, Berlin, Heidelberg, 2008.
15. Jurafsky, D., Martin, J., and Kehler, A. *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*. MIT Press, 2000.
16. Keffe. *Wiktionary:frequency lists for tv and movie scripts*, 2006. [http://en.wiktionary.org/wiki/Wiktionary:Frequency\\_lists](http://en.wiktionary.org/wiki/Wiktionary:Frequency_lists); accessed 10-Feb-2010.
17. Kurtoglu, T., Campbell, M., Gonzales, J., Bryant, C., and Stone, R. *Capturing empirically derived design knowledge for creating concep-*

- tual design configurations. In Proceedings of the ASME Design Engineering Technical Conferences And Computers In Engineering Conference. DETC2005-84405, Long Beach, CA (2005).
18. Langacker, R. An introduction to cognitive grammar. *Cognitive science* 10, 1 (1986), 1–40.
  19. Langacker, R. *Cognitive grammar: a basic introduction*. Oxford University Press, USA, 2008.
  20. Lawson, B. Schemata, gambits and precedent: some factors in design expertise. *Design Studies: Expertise in Design* 25, 5 (September 2004), 443–457.
  21. Martinetz, T. M., Berkshire, S. G., and Schulten, K. J. Neural gas network for vector quantization and its application to time-series prediction. *IEEE Transactions on Neural Networks* 4 (1993), 558–569.
  22. Moss, J., Cagan, J., and Kotovsky, K. Learning from design experience in an agent-based design system. *Research in Engineering Design* 15, 2 (2004), 77–92.
  23. Mukerjee, A., and Dabbeeru, M. M. The birth of symbols in design. In Proceedings of DETC'09 (September 9-12 2009), 2009 ASME Design Engineering Technical Conferences.
  24. Nanda, J., Thevenot, H., Simpson, T., Stone, R., Bohm, M., and Shooter, S. Product family design knowledge representation, aggregation, reuse, and analysis. *AI EDAM* 21, 02 (2007), 173–192.
  25. Park, S., and Gero, J. Qualitative representation and reasoning about shapes. In *Visual and Spatial Reasoning in Design* (Key Center of Design Computing and Cognition, University of Sydney, Sydney, Australia, 1999), vol. 99, pp. 55–68.
  26. Sarkar, S., Dong, A., and Gero, J. S. Design optimization problem reformulation using singular value decomposition. *Journal of Mechanical Design* 131, 8 (2009), 081006–1–10.
  27. Satish, G., and Mukerjee, A. Acquiring linguistic argument structure from multimodal input using attentive focus. In *Development and Learning, 2008. ICDL 2008. 7th IEEE International Conference on* (Aug. 2008), pp. 43–48.
  28. Schoen, D. A. Designing: Rules, types and words. *Design studies* 9, 3 (1988), 181–190.
  29. Steels, L. (2003). Evolving grounded communication for robots. *Trends in Cognitive Science*, 7(7), 308-312. doi:10.1016/S1364-6613(03)00129-3
  30. Yaner, P., and Goel, A. Analogical recognition of shape and structure in design drawings. *AI EDAM* 22, 02 (2008), 117–128.