

DETC2009-86709

THE BIRTH OF SYMBOLS IN DESIGN

Amitabha Mukerjee*

Computer Science & Engineering
Indian Institute of Technology Kanpur
Uttarpradesh, Kanpur, India
Email: amit@cse.iitk.ac.in

Madan Mohan Dabbeeru

Center for Robotics
Indian Institute of Technology Kanpur
Uttarpradesh, Kanpur, India
Email: mmadan@iitk.ac.in

DRAFT PAPER

ABSTRACT

In the widespread endeavour to standardize a vocabulary for design, the semantics for the terms, especially at the detailed levels, are often defined based on the exigencies of the implementation. In human usage, each symbol has a wide range of associations, and any attempt at definition will miss many of these, resulting in brittleness. Human flexibility in symbol usage is possible because our symbols are learned from a vast experience of the world. Here we propose the very first steps towards a process by which CAD systems may acquire symbols is by learning usage patterns or image schemas grounded on experience. Subsequently, more abstract symbols may be derived based on these grounded symbols, which thereby retain the flexibility inherent in a learning system.

In many design tasks, the “good designs” lie along regions that can be mapped to lower dimensional surfaces or manifolds, owing to latent interdependencies between the variables. These low-dimensional structures (sometimes called chunks) may constitute the intermediate step between the raw experience and the eventual symbol that arises after these patterns become stabilized through communication. In a multi-functional design scenario, we use a locally linear embedding (LLE) to discover these manifolds, which are compact descriptions for the space of “good designs”. We illustrate the approach with a simple 2-parameter latch-and-bolt design, and with a 8-parameter universal motor.

1 Efforts towards standardizing the design vocabulary

Evolving a standardized vocabulary for design has emerged as an important focus in engineering design with a need for communicating between differing systems and design groups. Possible applications include developing design repositories [Bohm et al., 2005, Nanda et al., 2007], computer assisted conceptual design [Gorti and Sriram, 1996, Campbell et al., 2000, Kurtoglu et al., 2005], [Chakrabarti et al., 2005, Gero and Fujii, 2000], etc. It is clear that vocabularies are structured, that is there are considerable relations between terms. Often, this is viewed as an ontology or as a structured relationship that captures a part of the semantics of these terms. One popular view of the engineering system considers the flow of energy, information, etc, and proceeds downward into detailed design. With its roots in value engineering ideas from the 1940s, these notions were seeded by the analysis in Pahl and Beitz [Pahl and Beitz, 1996] and a particularly influential study by Welch and Dixon [Richard and Dixon, 1994], leading to modern ontological models like the widely used *functional basis* model [Hirtz et al., 2002] or implementations on ontology tools [Nanda et al., 2007, Szykman et al., 2001].

1.1 The semantics of design symbols

All these models define a number of symbols at different levels in the hierarchy. Unfortunately the term “symbol”, as it is used in the logic and computational community is considerably different from its usage in cognitive linguistics and in everyday life. In the latter usage, symbols are imbued with meaning grounded on experience, whereas in the formal usage, it is merely

*Address all correspondence to this author.

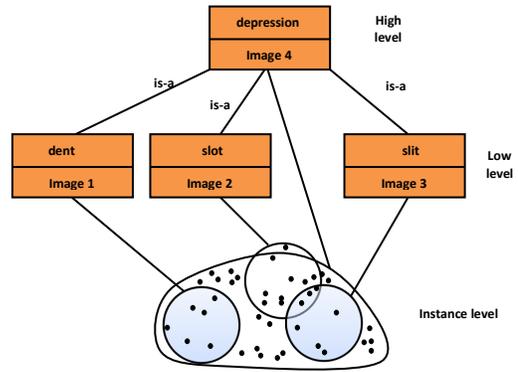
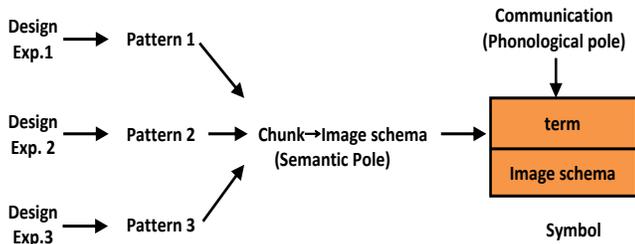


Figure 1. *Emergence of symbols based on experience*: Often the same abstract pattern (or *chunk*) appears in many experiences (e.g. the notion of “fit” for peg in hole, bolt in latch, plug in sink, etc.). If a chunk is valuable in compactly representing many situations, it has a higher likelihood of being communicated, thus acquiring a phonological pole and becoming a symbol. A symbol can then form other associations besides the initial chunk, all of which together constitute its semantic pole or *image schema*.

Figure 2. *Abstraction starts with ground instances*: Each symbol in this hierarchy has a term or label (“dent”, “slot” or “slit”) and a corresponding abstract pattern or “image schema”. The image schema is used in identifying an instance as belonging to a symbol category, but also in composing symbols, and in interpreting higher abstractions. Primitive design ontologies are born *is-a* through usage; when instances already known as dents or slots are also labelled as “depressions” by a trusted user, the system learns the subclass relationship. This makes grounded instances available even for the more abstract symbols. Similarly, other relations e.g. “dents are generally undesirable” would also be learned through usage and become part of the image schema. The number of such associations for each symbol is often very large, and limiting these to a few user-determined definitions is a major contributor to brittleness in knowledge systems.

a token constructed from some finite alphabet, and is related only to other terms. If we may present an analogy, the symbol “red” as it is used in a computer, is akin to the understanding of the symbol by a blind man; he knows that it is a “colour” (whatever that may be) and that “green” and “blue” are other colours, and maybe even that “crimson” and “vermillion” are shades of “red”, but his understanding of red is dramatically different from that of a sighted person, because the semantic pole is not connected to direct experience.

On the other hand, “symbol” has come to be understood in cognitive science (and also traditionally in linguistics, e.g. de Saussure ([De Saussure, 1986]), as the tight binding of the of the psychological impression of the sound (the “phonological pole”) with the mental image of the meaning (the *semantic pole*) [Langacker, 1986]. The mental image or image schema includes all sorts of associations and is somewhat different for each user, though social convention ensures a degree of overlap between mental images within the language community.

However, the notion of symbol is more far-reaching than communication. It turns out that to some extent, the symbols help divide up the world into chunks, and eventually, it may reflect changes in how we think. As an example relevant to the design community, we may mention how the Korean language makes a distinction between spatial tight-fit situations, *kkita*, (as in “put the cap on the pen”, “hand in glove”) from other usages of “in” or “on”. Infants growing up in English and Korean linguistic environments were sensitive to both contrasts, but English children appear to lose this sensitivity around the time they start acquiring language, suggesting that the language construct may have weakened their sensitivity to these changes [McDonough et al., 2003].

Symbols in CAD systems are not completely devoid of the semantic pole, indeed, implemented systems usually instantiate these symbols in terms of geometrical or physical attributes, but these details are largely implementation specific, heuristic, and are often not transferable to new domains. At the higher, or abstract levels in the design hierarchy, there may be some agreement on terms but as these get instantiated into specific instance, many definitions tend to be simplistic or even arbitrary. In many situations, the same physical and geometric symbol may require different instantiations in different domains, which is akin to the notion of *polysemy* in language (e.g. a “slot” on a coin-machine vs a “slot” on a machining fixture). These type of situations are almost impossible to program, since the range of usage very large, and may even be unbounded.

On the other hand, incompatibility of design vocabulary is rarely a problem between humans (that’s why exceptions often become memorable). If designers A and B are talking, and A does not have a particular symbol, its image-schema may emerge through a small amount of discussion; in many cases, just a single example may be enough to stretch an existing concept in A to the current one. Of course, this new symbol remains a bit wobbly, and designer A is aware of it, and subsequent uses of the symbol will serve to ground it. All this is possible because the semantic

pole for the human is a complex, elastic set of associations that cannot be defined in terms of a single predicate or even a range, it is the set of all situations where the symbol may be encountered (figure 2). All these associations need to be learned, and cannot be inferred based on a single definition (not to mention issues such as nonmonotonicity); hence the programmer-given single definition, usually created to demonstrate the example at hand, is a hopelessly inadequate semantics for a design symbol; and that is why we need bottom-up symbol discovery in order to ground a design vocabulary.

1.2 Bottom-Up Semantics in design

An alternative that has been proposed for modeling design concepts is to attempt to move more towards the human process, to learn symbols based on design experience [Gero and Fujii, 2000]. The human design process is a constant, motivated exploration of the design space, e.g. through sketching. All the while, the designer is focusing on the designs that are “good” in some functional sense, and eventually, some kinds of patterns emerge as the common characteristics of these designs. This is one sense in which sketches “talk back” to the designer [Goldschmidt, 2003]. These patterns result in constraints whereby many of the initial design variables can be combined, a process cognitively known as *chunking* [Gobet et al., 2001].

Considering the task of designing a padlock, we may learn that in a padlock, to balance the strength in its components, the shackle diameter increases roughly in proportion with body size. Thus these two parameters can then be brought down to a single chunk. Some other parameters also follow this trend, and also become part of this chunk. Thus, this enables the expert designer to consider a wider range of variables, than would have been possible initially. It is well known that a designer who is an expert in a particular design domain is “confident of immediately choosing a good [design] based on experience” [Gross, 1986], and this is partly possible because of these patterns that she has internalized. Often, these relations remain implicit, so that the human designer justifies the decision vaguely as “looks right” [Ahmed et al., 2003]. However, if these patterns do cross the consciousness boundary and get explicitly noticed, the designer may recognize it as an “aha!” moment.

An early attempt at discovering patterns in the design space of shapes may be seen in relation to 2D shapes in the work of [Park and Gero, 1999]. Another interesting approach to learning some of the semantics is found within the tradition of design knowledge ontologies [Moss et al., 2004]. Here a learning layer, operating as a manager (M-agents), are 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 trends. These can then be used to extract chunks which are added to memory, and become part of the variable space. Further, the effectiveness of a new chunk can be tracked in subsequent designs to ascertained

its utility.

Recently, an interesting approach to uncovering semantic connections that relate design variables with parameters has been proposed by [Sarkar et al., 2008], which considers the co-occurrence matrix of variables and constraints. Performing a Singular Value Decomposition on this matrix, one can use the ordered eigenvalues as an indication of the relevance of the relations between different variable groups. This provides a significant step in permitting designers to identify which variables to keep as design variables (as opposed to relegating it as a dependent parameter). By limiting one’s attention to the top few eigenvectors, one can also achieve a significant reduction in dimensionality.

1.3 Emergence of design concepts

None of these proposals however learn the concepts underlying the symbol (the semantic pole) in a grounded manner, and therefore lack the flexibility of the human designer. By *grounded*, we refer to the progressive manner in which a human designer learns her concepts - the more abstract ones are based on earlier, concrete concepts, but are still presented through instances. In the end, many concepts are grounded in terms of a number of experiential instances. For a human designer, this learning is not limited to the years of training as a designer, but must include all of her knowledge about the world, the so called commonsense knowledge. Thus, the fact that a fat peg will not go into a thin hole is part of her prior knowledge. Indeed, it is likely that the process by which she acquires these patterns, built upon many layers of pre-existing knowledge, may be similar in some salient ways with her earliest learning.

In this work, we propose to take the first step towards building this grounded semantics, which we call the *birth of symbols*. We adopt the cognitive interpretation, taking a symbol to mean the tight binding between a label and a large set of consciously accessible experiential patterns. The label is usually a communicative term, and constitutes its phonological pole. Occasionally, however, for example while “talking” to a sketch, a designer may get a conscious awareness of a constraint without verbalizing it - and these may also be a kind of symbol. Later, it may be verbalized, and acquire a name. The conscious awareness process (called *reification*) results in a new symbol. At the same time, amorphous implicit schemas, which are formed subconsciously (e.g. [Gladwell et al., 2005]) are incipient symbols, but not quite there yet. Still, they are useful, and we do use them, but they need to prove their mettle before they become true symbols. This interpretation is in line with a long tradition in psychology and linguistics [Mandler, 2004].

In the computational version of this learning process, we use several learning paradigms. One learning process, based on the training data of “good designs” is a simple general purpose function approximation - we use a single layer perceptron [Bishop,

2006]. Such function learners, while they can learn the patterns for “good designs”, are not capable of the information compression that is necessary in birthing symbols. There are several approaches that find such data encapsulations - when the clustering involves convex (linear) combinations of the variables, one may use PCA. Where the patterns sought are non-linear (imagine the spiral layer of jam in a swiss-roll), then these are better learned by a number of non-linear manifold learning techniques; here we use the well-known Locally Linear Embedding (LLE) [Saul and Roweis, 2003].

Since we are concerned with discovering patterns that apply to the set of *good* designs, we must have a way of assessing whether a design is good or not. This implies that some set of functional criteria must be available; even if it is not quantitative, it must be capable of ordering a wide variety of design instances, and given a design instances (a set of values for the design variables) it must be possible to determine the degree of functional feasibility. Such a model of function is usually not available in the very earliest stages of design; indeed even for human designers, insights usually arise while fiddling with instances, when suddenly a pattern may light up. So we are considering this type of emergence only after an embodiment is available, when the nature of its function can be related to the set of design variables. As in [Moss et al., 2004] we consider “good designs” as those arising from multi-criteria optimization or based on user-defined minimal functional criteria.

The next sections outline the development of the infant designer into a (very) early designer; the idea is that computational models capable of human-like ability would be possible based on the equivalent of years of exposure to such situations. In the end, we may expect to have a grounded symbol system that can reason about the flow of energy and take it down to a detailed design, but that end, even if it is achievable, is clearly far away. At present, let us start with the infant designer.

2 Infant designer

An infant designer is really the baby who is first discovering regularity of object behaviour in the world and is one who is just beginning to form her knowledge of the world. She can make various choices, and evaluate them based on some notion of function. Considering the peg-in-hole task alluded to, we see how she might learn the concept that a peg must be smaller than a hole.

The functional model considered is simple - the design is functionally feasible if the peg can go in (actually our system computes the configuration space - the penetration region disappears when $w < t$ - but this is not relevant here). We consider a horizontal version of the peg-in-hole - a latch is entering a slot on a bolt, say. Figure 3 shows how after evaluating a number of instances in the design space of latch-widths w and slot-widths t ; in (w, t) space, a clear 45 degree line emerges, separating the

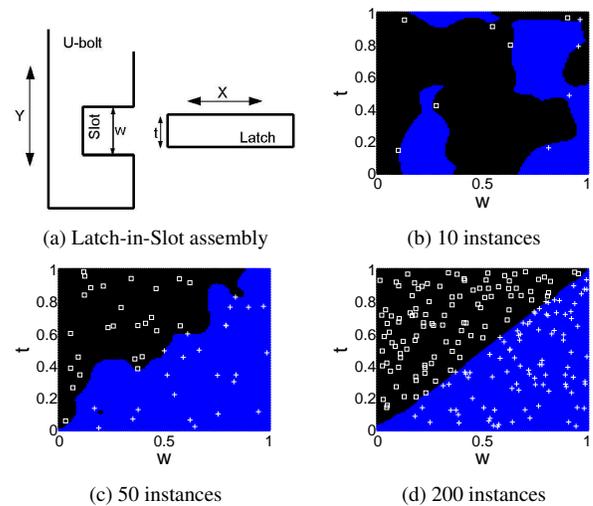


Figure 3. *Learning through experience that latch-must-be-smaller-than-slot ($w > t$).* (a) A latch of thickness t is fitted to a slot of width w . The learned patterns are shown in (w, t) -space in (b)-(d). The quality of the learned pattern varies greatly with degree of experience: results shown for a multi-layer perceptron after experiencing 10, 50, and 200 design instances.

“good designs” from the bad.

Does this constitute symbolic knowledge for the infant designer? Most likely not. However, it is something that might become a symbol as she acquires other concepts that she can refer to. An important aspect of symbols is that they are defined in relation to other symbols, e.g. if we consider vocabulary A which has both the terms “slot” and “slit”, versus system B, which has only “slot”, then we might expect the semantic pole of “slot” to more narrowly defined in A than in B. Similarly for the infant, it needs a certain density of putative symbols before making the symbolic jump.

What is interesting in the results of figure 3 is how, after experiencing just a few instances, the pattern is inchoate, so the baby keeps trying to insert fat pegs into the thin holes, filling up the negative (black) area of the figure. Eventually the defining boundary becomes sharper, and at some point it can be said to know the principle, at least implicitly.

At the next step for our infant designer, we consider the concept that a designer knows as “fit”. By now our infant learner will attempt to insert pegs only if they are smaller than the slot. The functional definition for the degree of *fit* may be: how much does it wiggle? Defining the wiggle in terms of the area of the free-space in the configuration space, we see that if the wiggle desired is very small, we get the situation on the left, and if it is very large, we get the situation on the right. Eventually, the learner learns the concept of “fit” as a chunk (composed as $w - t$) - thus, given a level of fit, it imposes a constraint where w and t

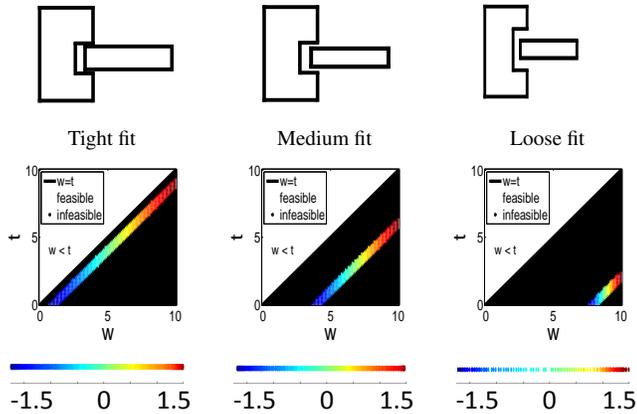


Figure 4. *Birth of the image-schema for “fit”*: An insertion task with different kinds of fit are shown in the top row and the corresponding design spaces (w, t) with feasible and infeasible regions are shown below. The function is given as the amount of play available (amount of free-motion or wiggle). If the desirable wiggle is specified, the two-dimensional design space is effectively reduced to one since a relation emerges between the feasible w and t . This mapping or image schema is a early prototype of the concept of “fit”.

are related in a manner where they constitute a one-dimensional chunk instead of two independent variable.

Of course, from a machine learning perspective, both these examples are rather elementary. Our objective in presenting it is merely to emphasize the role of even the earliest knowledge in many advanced design situations. These two concepts are also among our earliest knowledge achievements; typically, infants learn containment (peg in hole) by about 3 months, and tight vs loose by 5 months [Casasola et al., 2003]. Many cognitive scientists believe that our concepts of abstraction, including the *is-a* construct used in ontologies, is a metaphorical extension of containment [Lakoff and Johnson, 1999].

3 Symbol emergence

As the designer matures from infancy, we can consider the more general process by which symbols form. These correspond to the stages shown in figure 5. At first, the designer explores with instances in the design space, distinguishing the good designs from the bad. Eventually a region in the design space emerges as that containing mostly the feasible designs - this is the Functionally Feasible region (FFR), or the space of “good designs”. At least implicitly, this region is being continuously abstracted in terms of function. If the region is a very simple, one may even get an explicit characterization for it. However, in many high dimensional design spaces (even in the 8-dimensional example next) such relations are far from obvious.

The other interesting aspect of the FFRs is that they often

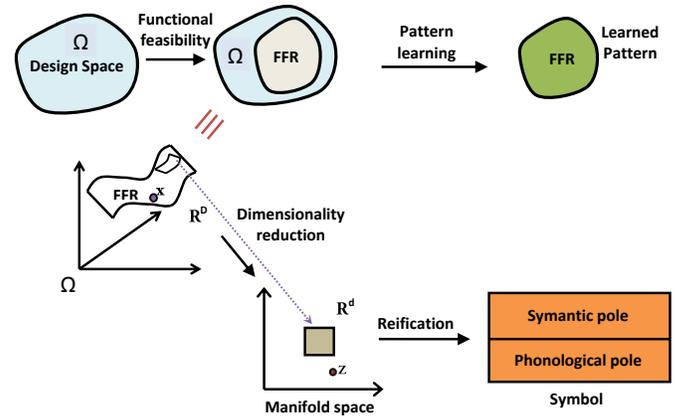


Figure 5. *The symbol emergence process*: our main interest is to discover and learn structural or behavioral chunks that result in good designs, corresponding to functionally feasible regions (FFRs) in the designs space. FFRs typically reflect multiple functional criteria, and may be obtained from some approximate optimization, or from user specified minimal functional criteria. A set of FFR instances can be used to learn a pattern of functional feasibility, the quality of this pattern improves with experience as earlier. Once the FFR is sufficiently rich, one may also discover that they lie along some low-dimensional manifold (R^d) embedded in the high-dimensional design space R^D ($d \ll D$). The lower dimensional space is then a chunked representation for the initial design space. If this relation becomes conscious, it may then become a design symbol.

correspond to narrow bands of functional feasibility. This may be because they are the result of (possibly subconscious) multi-objective optimization - thus, if there are k design objectives, then they constitute a $k - 1$ surface in the objective space. When the function measures that map from the design variable space to the objective space are at least continuous, their Jacobians would be well-posed, and the near neighbours in the objective space may correspond to near neighbours in the design space. If this holds, we may expect the designs to lie along a $k - 1$ surface (or “manifold”) in the design space as well (shown as a folded patch in the figure). More generally, the mapping from the design space to the functional objective space is not so well-posed, but the dimensionality reduction, while not quite as pronounced, may still be very significant.

Each dimension in this reduced dimensional map reflects an inter-relation between many independent design parameters. Subsequently, many other design situations also reveal such reduced dimensional hyper-surfaces, with different dimensionalities. Some of these reduced-dimensional mappings or chunks may recur in other situations - this makes the chunk useful, which

is an important criteria for becoming a symbol. In the interim, the designer may use these chunks with a dim awareness of it for a long period, even several years. Meanwhile, as other instances are being explored, and the confidence in the manifold mapping increases, the designer may eventually articulate it (at least to herself), which would be the birth of a true symbolic representation for this concept. At this point, a label may get attached to it, and many other associations would eventually accrue to this term / image-schema pair; it would then constitute a truly reified symbol in the cognitive sense.

In our computational analogue of this process, a key step is that of discovering the lower-dimensional patterns; this is the algorithm we describe next. This algorithm will then be applied to a well-known design problem, to obtain chunks.

3.1 Dimensionality Reduction

We now present the algorithm used for obtaining a low-dimensional representation for the “good design” subspace of the original high-dimensional design space. As noted above, the number of functional objectives is usually much less than the number of parameters in the design space. For example, a digital camera (in terms of all its components, and assembly processes) may have several hundred parameters, but only about ten functional measures. Based on these functional measures and the constraints among the design parameters, the good designs may be constrained to a much lower dimensional manifold in the design space. Thus, although the design is defined in terms of a hundred parameters, for the class of “good designs”, there are often many interrelations between these; each such interrelation constitutes a chunk or a dimension in the resulting low-dimensional surface or manifold.

Dimensionality reduction techniques attempts to find low-dimensional structures in high-dimensional space, using a large variety of approaches. There have been linear dimensionality reduction methods [Bishop, 2006] - e.g. independent component analysis, linear discriminate analysis, principal component analysis. These linear methods fail when the data lies on a non-linear manifold; in such situations the linear algorithms give the smallest convex subspace encapsulating the manifold, which is often of a much higher dimension. In practice, non-linear relations between design variables are extremely common, and in such situations, nonlinear dimensionality reduction yields superior results [Tenenbaum et al., 2000]. Approaches for obtaining the non-linear representation of the data include global methods (Isomaps [Tenenbaum et al., 2000], Laplacian Eigenmaps [Belkin and Niyogi, 2002]) and local methods (Locally Linear Embedding [Saul and Roweis, 2003]). Local approaches try to preserve the local geometry of the data. By approximating each point on the manifold with a linear combination of its neighbors, and then using the same weights to compute a low-dimensional embedding, LLE tries to map the nearby points on the manifold

to nearby points in the low-dimensional representation.

Algorithm 1 Local Linear Embedding [Saul and Roweis, 2003]

1. Compute the neighbors X_j of each data point, X_i .
 2. Compute the weights W_{ij} that best reconstruct each data point X_i from its neighbors, minimizing the reconstruction error ($\epsilon(W) = \sum_i |X_i - \sum_j W_{ij} X_j|^2$) by constrained linear fits.
 3. Compute the vectors Γ_i best reconstructed by the weights W_{ij} , minimizing the quadratic form ($\Phi(\Gamma) = \sum_i |\Gamma_i - \sum_j W_{ij} \Gamma_j|^2$) by its bottom nonzero eigenvectors.
-

The Locally linear Embedding (LLE) is an eigenvector method for the problem of nonlinear dimensionality reduction (algorithm 1. [Saul and Roweis, 2003, Roweis and Saul, 2000]). The assumption underlying LLE is that the manifold embedded in a high-dimensional space is locally linear if each data point and its neighbors lie approximately on a locally linear patch on the manifold (the manifold considered to be hyper-flat locally). Thus each data point can be approximated as a weighted linear combination of its nearest neighbors. The weights or coefficients of this approximation characterize the local geometries in a high-dimensional space. The key insight is that for a true manifold, the same weights would also apply in the low-dimensional mapping. This can then be used to find low-dimensional embeddings preserving the relative weights in its neighbourhood. In this study, we use LLE to identify the underlying non-linear manifold in high-dimensional design space.

3.2 Universal Motor example

Next we consider a well-known problem in the design literature, that of designing an Universal Motor. The Universal Motor is a motor with a wound armature and a wound stator. The armature is fed via brushes on a commutator, and is essentially the same as a D.C. motor. The universal motor operates off a single phase A.C. supply and accelerates until the load torque equals the output torque. Such motor are used from home appliances like domestic vacuum cleaners, food processors, mixer grinders to industry applications like material handling, power sector motors. Universal motors have been well studied in the product family design literature [Simpson, 1998], where the physical description and schematic of the universal motor have been presented. The design space for embodiment design consists of 10 design variables $\vec{v} = \{N_c, N_s, A_{wa}, A_{wf}, r_o, t, l_{gap}, I, V_t, L\}$ where N_c : number of wire turns on the armature, N_s : number of wire turns on each pole on the field, A_{wa} : cross-sectional area of the wire on the armature, A_{wf} : cross-sectional area of the wire on the field, r_o : radius of the motor, t : thickness of the stator, I : current drawn by the motor, L : stack length, and the performance

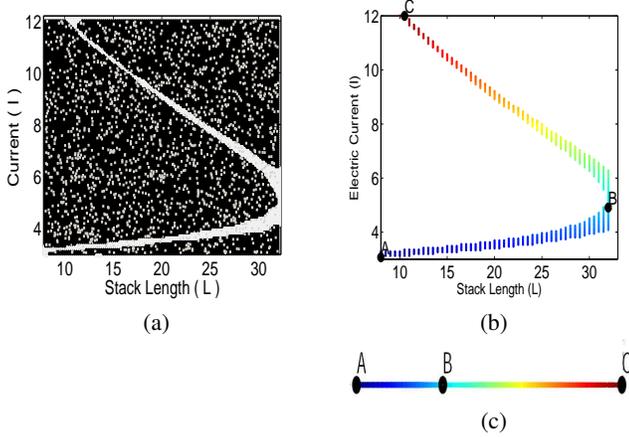


Figure 6. *Chunking on the L, I subspace for Universal Motors:* (a) The implicit constraint on the L, I subspace of the Design Space is learned for 2000 design instances under the functional specification $280 \text{ W} < \pi_{power} < 295 \text{ W}$. (b) The feasible designs in the L, I subspace, in which each design instance can be identified based on two parameters L and I . This subspace is mapped into a low-dimensional (one-dimension) shown in (c). A, B, and C are three different design instances mapped from L, I space to one-dimensional space.

behaviors are taken as strength, mass, energy and efficiency and the corresponding performance metrics in terms of these design variables can be $\pi_{torque}(\vec{v}) = \frac{N_c \Phi I}{\Pi}$, $\pi_{mass}(\vec{v}) = mass_{windings} + mass_{armature} + mass_{stator}$, $\pi_{power}(\vec{v}) = V_t I - I^2 (R_a + R_s) - 2I$, and $\pi_{efficiency}(\vec{v}) = \frac{\pi_{power}}{V_t I}$. The derivation and the related equations are presented well in [Simpson, 1998]. And the main constraints we considered here for the feasible motors are (i) the magnetizing intensity $H < 5000$ and (ii) the outer radius of the stator r_o must be greater than the thickness of the stator t .

We first proceed with the same learning process as earlier, where we are abstracting the design experience in terms of patterns of FFRs. These learned patterns can be used in future design tasks of similar designs. A minimal parameter set for the universal motors may be considered in terms of the two design parameters L and I , while keeping other parameters constant [Simpson, 1998]. To obtain the feasible universal motors for the desired functional range of power $280 \text{ W} < \pi_{power} < 295$. Figure 6(a) shows the result of learning the FFR (the valid designs resulting from this constraint). These lie along a small band, which can be thought of as a curved 1-D manifold (with a slight thickness). 6(b).

Next, we demonstrate the manifold nature of the FFR using the eigenvector method called locally linear embedding or *LLE*, explained in the next section. The mapping between the nonlinear feasible region (Fig. 6 (b)) and the one-dimensional chunk for it below (Fig. 6 (c)) shows the continuity of mapping between these. If we take three data points A, B, and C in L, I

	X		Γ	π_{torque} (Nm)	π_{mass} (kg)	π_{power} (W)	$\pi_{efficiency}$
	L(m)	I (Amp)					
A	32.0	4.0909	-0.2102	0.2643	0.8428	280.4795	0.5962
B	22.5	3.5455	-0.1430	0.6815	0.7613	289.1254	0.3580
C	10.5	12.000	0.0077	0.7462	0.4048	282.8583	0.2050

Table 1. *Lower dimensional FFR:* To map these FFRs in L, I space to Γ space we have used local linear embedding algorithm (shown Algorithm 1) with three nearest neighbors ($K=3$) for each data point in X space. All data points are mapped from $X_{D \times N}$ to $\Gamma_{d \times N}$, where D is the dimension of the L, I space ($D = 2$), d is one-dimensional Γ space and N is the number of data points. Three data points A, B and C are showing mapping from L, I space to Γ space.

space. Let us say $X = [A \ B \ C]$, each data point is a real-valued vector, with of dimensionality 2. With the help of Local Linear Embedding (LLE) algorithm [Roweis and Saul, 2000], we construct a neighborhood preserving mapping from L, I space to Γ . The particulars of these three points mapping is shown in Table 1.

After this γ chunk is discovered in this Universal motor situation, it is possible that it may also emerge in other situations. If so, a designer may communicate the idea of this chunk to another; let us say he calls it “gavagai”. Then as the term “gavagai” spreads in the design community, it would occur in many other situations, and each association would form part of the semantics of the term gavagai.

Indeed, different generations of designers would discover their own image schemas for gavagai, and the term may be applied to wider contexts (“generalization”) or more specific ones (“specialization”); thus, it would slowly change its semantics over time. This is another way in which static programmed machine semantics, even if they can capture all the usages at a given point of time, would not be able to keep up with human usage unless it adopts an experiential learning modality as proposed here.

4 Chunking from Multi-Objective Optimization

Until now, we have considered the FFRs as arising from user defined functional desiderata. However, how does the user arrive at these functional bounds? In practice, a more useful approach towards finding FFRs is to consider them as the estimated pareto fronts arising from Multi-objective optimization. This reflects, in a very approximate manner, the human designer’s search for optimality, or at least some degree of optimality, in complex design spaces.

Here, we need not know the bounds on the performance, just the set of performance measures, computable given an instance of the design, are to be specified. In general, we consider the

design problem as one of balancing multiple objectives, some of which may be conflicting. If a solution A is better than solution B in all the functional criteria, we say that A *dominates*— B. If a solution is such that no other solution dominates it, then it is a possible optimum, the set of all such non-dominated solutions (the *pareto-front*), usually lies along a surface in the space of objective functions; if there are k functions and these are independent, then the dimensionality of this surface is $k - 1$. Most computational multi-objective optimization (MOO) algorithms provide estimates for the non-dominated front.

For our purposes, the non-dominated front constitutes a description of “good designs”. The front is obtained in the space of objective functions, and our first task would be to map it to the design space. We would then consider these “good designs” in the design space and try to find a manifold representation for it.

4.1 Example A : Universal Electric Motors

Let us now formulate the design of Universal motors as a multi-objective optimization problem to identify FFRs. Here we consider eight design parameters $\vec{v} = \{N_c, N_s, A_{wa}, A_{wf}, r_o, t, I, L\}$ and their ranges are, N_c : number of wire turns on the armature [100, 1500], N_s : number of wire turns on each pole on the field [1, 500], A_{wa} : cross-sectional area of the wire on the armature [0.01, 1.0mm²], A_{wf} : cross-sectional area of the wire on the field [0.01, 1.0]mm², r_o :radius of the motor [1.0, 10.0]cm, t : thickness of the stator [0.5, 10.0]mm, I : current drawn by the motor [0.1, 6.0]A, L : stack length [0.057, 5.18]cm. The mathematical formulation of multi-objective optimization problem as follows:

Multi-Objective Optimization

$$\begin{aligned}
 &\text{Minimize} && \pi_{mass}(\underline{v}) \\
 &\text{Maximize} && \pi_{efficiency}(\underline{v}) \\
 &\text{Maximize} && \pi_{torque}(\underline{v}) \\
 &\text{Subject to} && g_1(\underline{v}) \equiv r - t > 0 \\
 & && g_2(\underline{v}) \equiv 5000 - H > 0, \\
 & && g_3(\underline{v}) \equiv 2.0 - \pi_{Mass} \geq 0, \\
 & && g_4(\underline{v}) \equiv 0.5 \leq \pi_{torque} \leq 5.0, \\
 & && g_5(\underline{v}) \equiv 300 \leq \pi_{Power} \leq 600 \\
 & && g_6(\underline{v}) \equiv \pi_{efficiency} - 0.15 \geq 0
 \end{aligned} \tag{1}$$

The multi-objective optimization (MOO) technique used here is NSGA-II [Deb, 2001], a well known evolutionary algorithm. The NSGA-II parameters used in this study are as follows: population size = 2000, maximum generations = 500, probability of crossover for both real-valued and binary variables = 0.8, probability of mutation 0.33 and 0.1 for real and binary variables and the distribution index for crossover and mutation are 16 and 30. The Pareto optimal front obtained using NSGA-II is shown in Fig. 7(a). The Pareto-optimal surface is for maximizing

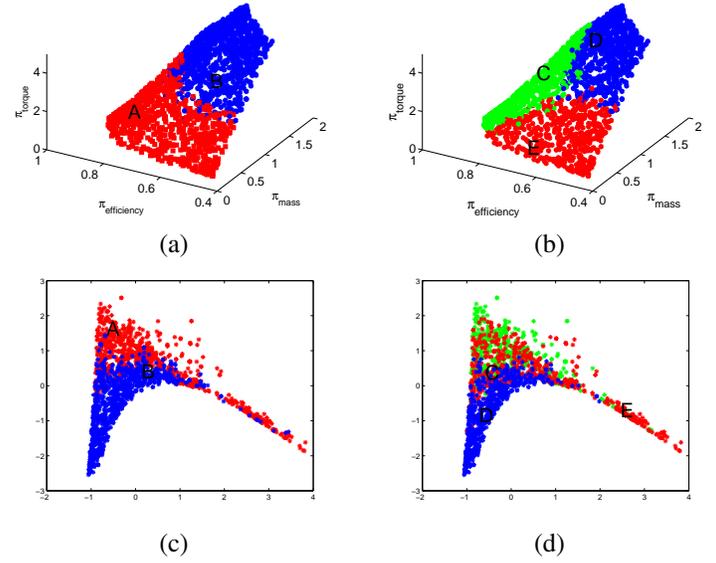


Figure 7. *Clusters in the non-dominated space for Universal motor.* (a) The non-dominated solutions (pareto-front) in the 3-objective space of mass, efficiency and torque. By considering the feasible designs in the design space, we obtain two clusters based on unsupervised clustering in the design space. (c) The manifold space corresponding to the map from the high-dimensional design space $D = 8$ to low-dimensional design space $d = 2$ obtained with the help of LLE. Similarity three clusters are formed with varying unsupervised clustering input parameters (b) and the corresponding low-dimensional manifold is shown in (d).

both the torque (π_{torque}) and efficiency ($\pi_{efficiency}$) while minimizing the mass (π_{mass}). Having obtained non-dominated sets of designs, and mapping these to the design space reveals that the good designs (FFRs) are often restricted to a few patches on a low-dimensional manifold, thus resulting in significant dimensionality reductions for the design space.

5 Intrinsic Dimension

The main parameters in this reduction algorithm is estimating (i) number of neighborhoods (K) for each data point in the higher dimension (D), and (ii) the lower dimensionality (d). The estimated dimension d of d^* is provided by the user as a parameter to the algorithm. If d is an under estimate of d^* , there is a loss of information and if d is an over estimate of d^* , then LLE will include arbitrary dimension. In PCA, the dimensionality is estimated by the number of eigen values of the simple covariance matrix comparable in magnitude to the largest eigen value. In LLE, to estimate the intrinsic dimensionality (d^*), one may count the number of eigen values of the covariance matrix ($M = (I - W)^T(I - W)$) (see [Saul and Roweis, 2003]) compar-

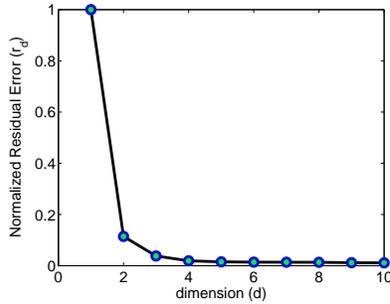


Figure 8. *Dimensionality of manifold for Universal Motors based on* . The FFR data is mapped onto manifolds of different dimensions, and then mapped back to the original design space and the error is estimated. The error drops sharply from 1-D to 2-D manifold, and then less sharply. The knee of the curve at “2” is indicative of the intrinsic dimensionality of the space. A separate maximum likelihood method estimates the dimensionality of the Universal Motor space to be 2.6.

ble in magnitude to the lowest non-zero eigen values [Polito and Perona, 2002]. But these eigen values are not informative based on the work [Saul and Roweis, 2003].

To obtain an estimate of the dimension of the manifold for our data set, we use the technique based on the idea that a dimensionality reduction algorithm should preserve information on a global scale, as measured using bijection [Martin and Backer, 2005]. For a given input dataset $X = \{X_1, \dots, X_N\} \subset R^D$, the dimensional reduction algorithm such as LLE provide a reduced dimensional representation $Y = \{Y_1, \dots, Y_N\} \subset R^d$ of the original data set X . It is assumed that X lies on a manifold M embedded in R^D with intrinsic dimensionality d^* . The estimated dimension d of d^* is a parameter to the algorithm. As per [Martin and Backer, 2005], a good low dimensional representation of X can be expected only if we should be able to go back and forth from X and Y with no loss of information. Let us say, this algorithm provides a map $f : X \rightarrow Y$, mathematically an inverse function $f^{-1} : Y \rightarrow X$ exists such that $f^{-1}(f(X_i)) = X_i$ for all i and exists only if the estimated dimension $d \geq d^*$ of the manifold. To determine the existence of f^{-1} , we have considered the same measure as proposed by [Martin and Backer, 2005], squared residual error $(r_d) = \sum_i \|f_d^{-1}(f_d(X_i - X_i))\|^2$, where $f_d : X \rightarrow Y$ is map produced by LLE, which depends on the estimate d of the intrinsic dimensionality and $f_d^{-1} = Y \rightarrow X$ is a proposed inverse to f_d . To estimate the actual dimension d^* of the manifold M , we found r_d for different values of d . This residual error should be zero when $d \geq d^*$ so that we may determine d^* by finding the smallest value of d such that $r_d = 0$. Hence we can predict the intrinsic dimension by minimizing both r_d and d . By observing the behavior of r_d for different values of d shown in Fig. 8 we can predict the intrinsic dimension can be 2.

Also, we have used a method proposed by [Levina and

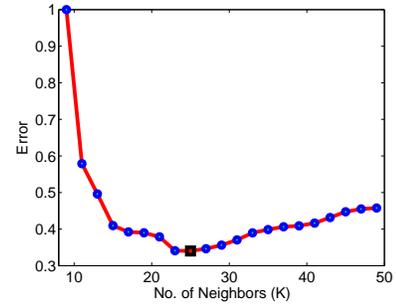


Figure 9. *Number of Neighborhoods:*

Bickel, 2004] for estimating the intrinsic dimension of a dataset by applying the principle of maximum likelihood to the distances between the neighbors. With this method we have obtained the intrinsic dimension value $d^* = 2.6$.

6 Selecting 'K' value

In LLE, the assumption is that local patches are linear, it means each of them can be approximated by a linear hyper-plane, and each data point can be represented by a weighted linear combination of its K nearest neighbors. The selection of number of nearest neighbors is an important parameter in LLE. A large number of nearest neighbors causes smoothing or eliminating the of small-scale structures in the manifold. A small number can falsely divide the continuous manifold into disjoint sub-manifolds [Kouropyteva et al., 2002]. We have experimented the hierarchical method proposed in [Kouropyteva et al., 2002] with our design data to select the optimum K value. In this process, without going through all steps of LLE, we can calculate the reconstruction error $\epsilon = \sum_{i=1}^N \|X_i - \sum_{j=1}^K W_{ij} X_{ij}\|^2$ for each K value varying from 1 to K_{max} . When varying K , the ϵ can be considered as function of K and return K is function of the coefficients (weights W_{ij}) and hence these coefficients alter as K changes. As a result, we can choose the K value which has the lowest ϵ . Fig. 9 is showing the $\epsilon(K)$ for our motor data and from this we found the optimal K value is 25.

7 Conclusion

We have presented an approach to the learning of design symbols which are understood to be a tight binding between a term or label and a semantic representation or image schema. The latter is an extensive set of experiential associations. These term-meaning pairs constitute an important part of the cognitive repertoire of the human designer, but the semantic pole is completely lacking or severely bleached in current machine systems. This makes it impossible to refer to instances of abstract symbols (defined in terms of other symbols) which is what gives the

human systems the plasticity in deploying these symbols.

In language learning by humans, only a few hundred symbols are learned *ab initio*, in the pre-linguistic stage. These are learned not just as the labels, but in terms of the rich set of associations that license its use. The remaining tens of thousands of lexical units in the adult vocabulary are learned by exposure to language, typically through reading. What this means is that symbols are ultimately defined in terms of other symbols, but even for these later symbols, instances are available as defined by the linguistic context. Symbol combinations are understood through a process of *composition* where the syntactic structure involving different symbols are combined to construct a more complex symbol [Langacker, 1986].

This process is clearly at work in the design process as well. Once the infant designer has learned a few symbols, these may be reinforced by communication, and be used to define other symbols. For example, it may be told that a “peg” is the thing that fits tightly in a “hole”. However, the symbol “fit-tightly-in-a-hole” cannot emerge unless one has experienced the “wiggle” associated with insertion tasks. Thus grounded or experienced symbols provide the plasticity needed in real-life applications; this is why human-defined symbols as used in design ontologies today are subject to brittle failures. In design terms, “syntax” may constitute the elements of a design and the modalities of connecting them, its semantics then is the overall behaviour and associations of the assembly. Thus, the semantics of fitting an electric plug into a socket may inherit some aspects of the symbol “fit-tightly-in-a-hole” but also some aspects of “energy-flow-electrical-connection”. The appearance of the typical plug, its prong structure, its weight, the sound of inserting it, all these would also be part of the semantics. Clearly, defining all of these experiential aspects would be a taxing task. Furthermore, with the evolution of technology (or language), the semantics may change, and maintaining such systems would be eventually a doomed enterprise. A grounded learning system is much more likely to be able to adapt to such changes.

This work is clearly exploratory, and much work is needed to define the conditions in which manifolds may exist, and the conditions under which a chunk may graduate into a symbol. However, we feel that non-linear manifold learning as an important step in discovering latent relationships among the many design parameters that permits defining this objective in terms of chunks, is a computational tool with enormous potential. An important question this leads to is how to compute the composition of more than one symbol; i.e. given the design elements each as an individual symbol, we need to be able to say what the conjunction of these elements (the syntax) will do, and whether the resulting object - a design instance - will be adequate to meet a design task or not. Again, depending on the “good designs” that emerge in the process, a combination of symbols may come to be designated as a symbol on its own right, leading to the birth of abstract symbols.

At the same time, there is an important role for high-level symbol ontologies. Just as designers can be taught certain relations, even computer functional models may benefit from explicit awareness of pre-programmed relations, especially at levels of abstraction that would be hard to grasp otherwise. However, it is essential that the semantic poles for these instructions be defined in terms of its own internal image schemas; without this, brittleness may again prevail. Further, shared ontologies are crucial to being able to communicate with humans.

In the final analysis, the argument presented here implies that in the long run, to create viable computer vocabularies for design, we must train the systems, at least in the earliest stages, to learn these relationships, by experiencing many design and real world situations. This may be done in an accelerated manner, but the system must be exposed to something like the vast array of experiences of a human - or possibly many more, since the abstraction algorithms available today may not be as efficient. As different systems are deployed in solving different problems, their somewhat differing input sets would result in somewhat different abstractions for the same symbols. These resulting design agents may therefore be somewhat less predictable than current computers, but such is the price of flexibility.

REFERENCES

- [Ahmed et al., 2003] Ahmed, S., Wallace, K. M., and Blessing, L. T. (2003). Understanding the differences between how novice and experienced designers approach design tasks. *Research in Engineering Design*, 14(1):1–11.
- [Belkin and Niyogi, 2002] Belkin, M. and Niyogi, P. (2002). Laplacian Eigenmaps and Spectral Techniques for Embedding and Clustering. *Advances in Neural Information Processing Systems*, 1:585–592.
- [Bishop, 2006] Bishop, C. (2006). *Pattern recognition and machine learning*. Springer.
- [Bohm et al., 2005] Bohm, M., Stone, R., and Szykman, S. (2005). Enhancing Virtual Product Representations for Advanced Design Repository Systems. *Journal of Computing and Information Science in Engineering*, 5:360.
- [Campbell et al., 2000] Campbell, M., Cagan, J., and Kotovsky, K. (2000). Agent-based synthesis of electromechanical design configurations. *Journal of Mechanical Design*, 122:61.
- [Casasola et al., 2003] Casasola, M., Cohen, L. B., and Chiarello, E. (2003). Six-month-old infants’ categorization of containment spatial relations. *Child Development*, 74:679–693.
- [Chakrabarti et al., 2005] Chakrabarti, A., Sarkar, P., Leelavathamma, B., and Nataraju, B. S. (2005). A functional representation for aiding biomimetic and artificial inspiration of new ideas. *AI EDAM*, 19(02):113–132.
- [De Saussure, 1986] De Saussure, F. (1916/1986). *Course in general linguistics*. Open Court.

- [Deb, 2001] Deb, K. (2001). *Multi-Objective imization using Evolutionary Algorithms*. Chichester, John Wiley and Sons, Ltd., 1 edition.
- [Gero and Fujii, 2000] Gero, J. and Fujii, H. (2000). A computational framework for concept formation for a situated design agent. *Knowledge-Based Systems*, 13(6):361–368.
- [Gladwell et al., 2005] Gladwell, M., Finder, H., and network, C.-S. T. (2005). *Blink: The power of thinking without thinking*. Penguin Books.
- [Gobet et al., 2001] Gobet, F., Lane, P., Croker, S., Cheng, P., Jones, G., Oliver, I., and Pine, J. (2001). Chunking mechanisms in human learning. *Trends in Cognitive Sciences*, 5(6):236–243.
- [Goldschmidt, 2003] Goldschmidt, G. (2003). The backtalk of self-generated sketches. *Design Issues*, 19(1):72–88.
- [Gorti and Sriram, 1996] Gorti, S. R. and Sriram, R. D. (1996). From symbol to form: a framework for conceptual design. *Computer-aided design*, 28(11):853–870.
- [Gross, 1986] Gross, M. D. (1986). *Design as Exploring Constraints*. PhD thesis, Department of Architecture, Massachusetts Institute of Technology.
- [Hirtz et al., 2002] Hirtz, J., Stone, R., McAdams, D., Szykman, S., and Wood, K. (2002). A functional basis for engineering design: Reconciling and evolving previous efforts. *Research in Engineering Design*, 13(2):65–82.
- [Kouropteva et al., 2002] Kouropteva, O., Okun, O., and Pietikainen, M. (2002). Selection of the optimal parameter value for the locally linear embedding algorithm. In *Proc. of the 1st Int. Conf. on Fuzzy Systems and Knowledge Discovery, Singapore*, pages 359–363.
- [Kurtoglu et al., 2005] Kurtoglu, T., Campbell, M., Gonzales, J., Bryant, C., and Stone, R. (2005). Capturing empirically derived design knowledge for creating conceptual design configurations. In *Proceedings of the ASME Design Engineering Technical Conferences And Computers In Engineering Conference. DETC2005-84405, in review, Long Beach, CA*.
- [Lakoff and Johnson, 1999] Lakoff, G. and Johnson, M. (1999). *Philosophy in the Flesh: The embodied mind and its challenge to Western thought*. Basic Books, New York.
- [Langacker, 1986] Langacker, R. W. (1986). An introduction to cognitive grammar. *Cognitive Science*, (10):1–40.
- [Levina and Bickel, 2004] Levina, E. and Bickel, P. (2004). Maximum likelihood estimation of intrinsic dimension. *Ann Arbor MI*, 48109:1092.
- [Mandler, 2004] Mandler, J. M. (2004). *Foundations of Mind : Origins of conceptual thought*. Oxford University Press, New York.
- [Martin and Backer, 2005] Martin, S. and Backer, A. (2005). Estimating manifold dimension by inversion error. In *ACM Symposium on Applied Computing*, pages 22–26.
- [McDonough et al., 2003] McDonough, L., Choi, S., and Mandler, J. (2003). Understanding spatial relations: Flexible infants, lexical adults. *Cognitive Psychology*, 46(3):229–259.
- [Moss et al., 2004] Moss, J., Cagan, J., and Kotovsky, K. (2004). Learning from design experience in an agent-based design system. *Research in Engineering Design*, 15(2):77–92.
- [Nanda et al., 2007] Nanda, J., Thevenot, H., Simpson, T., Stone, R., Bohm, M., and Shooter, S. (2007). Product family design knowledge representation, aggregation, reuse, and analysis. *AI EDAM*, 21(02):173–192.
- [Pahl and Beitz, 1996] Pahl, G. and Beitz, W. (1996). *Engineering Design: A Systematic Approach*. Springer.
- [Park and Gero, 1999] Park, S. and Gero, J. (1999). Qualitative representation and reasoning about shapes. In *Visual and Spatial Reasoning in Design*, volume 99, pages 55–68.
- [Polito and Perona, 2002] Polito, M. and Perona, P. (2002). Grouping and Dimensionality Reduction by Locally Linear Embedding. *Advances in Neural Information Processing Systems*, 2:1255–1262.
- [Richard and Dixon, 1994] Richard, W. V. and Dixon, J. R. (1994). Guiding conceptual design through behavioural reasoning. *Research in Engineering*, 6:169–188.
- [Roweis and Saul, 2000] Roweis, S. and Saul, L. (2000). Non-linear Dimensionality Reduction by Locally Linear Embedding.
- [Sarkar et al., 2008] Sarkar, S., Dong, A., and Gero, J. S. (2008). A learning and inference mechanism for design optimization problem (re)-formulation using singular value decomposition. In *Proceedings of DETC'08, ASME Design Engineering Technical Conferences and Computers and Information in Engineering Conference*.
- [Saul and Roweis, 2003] Saul, L. K. and Roweis, S. T. (2003). Think globally, fit locally: unsupervised learning of low dimensional manifolds. *The Journal of Machine Learning Research*, 4:119–155.
- [Simpson, 1998] Simpson, T. W. (1998). *A Concept Exploration Method for Product Family Design*. PhD thesis, Georgia Tech University, Dept Mechanical Engineering.
- [Szykman et al., 2001] Szykman, S., Sriram, R., and Regli, W. (2001). The role of knowledge in next-generation product development systems. *Journal of computing and information Science in Engineering*, 1:3.
- [Tenenbaum et al., 2000] Tenenbaum, J., Silva, V., and Langford, J. (2000). A global geometric framework for nonlinear dimensionality reduction.