

Algorithmic perception of vertices in sketched drawings of polyhedral shapes

Pedro Company¹, Raquel Plumed², Peter A.C. Varley², Jorge D. Camba³

¹ Institute of New Imaging Technology, Universitat Jaume I, Spain

² Department of Mechanical Engineering and Construction, Universitat Jaume I, Spain

³ Department of Computer Graphics Technology, Purdue University, USA

Abstract

In this paper, visual perception principles were used to build an artificial perception model aimed at developing an algorithm for detecting junctions in line drawings of polyhedral objects that are vectorized from hand-drawn sketches. The detection is performed in 2D, before any 3D model is available and minimal information about the shape depicted by the sketch is used.

The goal of this approach is to not only detect junctions in careful sketches created by skilled engineers and designers, but also detect junctions when skilled people draw casually to quickly convey rough ideas. Current approaches for extracting junctions from digital images are mostly incomplete, as they simply merge endpoints that are near each other, thus ignoring the fact that different vertices may be represented by different (but close) junctions and that the endpoints of lines that depict edges that share a common vertex may not necessarily be close to each other, particularly in quickly sketched drawings. We describe and validate a new algorithm that uses these perceptual findings to merge tips of line segments into 2D junctions that are assumed to depict 3D vertices.

Index Terms: Algorithmic perception. Polyhedral shapes. Vertices. Junctions.

1. Introduction

Engineers and designers regularly create sketches during the conceptual design stages of new product development processes [Jen92], [Pac05]. Designers rightly assume that engineers will be able to interpret such sketches, but current geometry-based algorithms cannot interpret sketches the same way and as efficiently as humans can.

Our long-term goal is to develop automatic interpretation capabilities for sketched drawings to subsequently reconstruct the corresponding 3D models. This problem can be divided into a number of sub-problems. The first is vectorization, which is a reconstruction stage that converts strokes into lines. Our hypotheses regarding vectorization are based on traditional perception principles, such as Gestalt Laws and Hoffman's rules, which we reformulated in an algorithmic fashion.

A particular stage in vectorization is the detection of junctions. Many computer-based reconstruction approaches assume that junctions are carefully depicted in the sketch ([XCS14] is a recent example). However, junctions are neither necessarily contained explicitly in the original sketch, nor automatically obtained after vectorization. Our general assumptions are that (a) the problem of detecting junctions in vectorized sketches is not merely geometrical but also perceptual (see [Kan79], cited by [JGH09]), and (b) that skilled people may draw both carefully and casually.

Sketches convey information that can be perceived by humans. Skilled engineers and designers often use two distinct types of sketches. Careful sketches are used to convey detailed and/or finished ideas, whereas casual sketches convey quick and/or incomplete ideas. The quality of the sketches of an average population can be represented as a continuum including “good” and “bad” sketches. This continuum is clearly biased towards good sketches for engineers and designers, which are trained sketchers. However, we distinguish between careful and casual sketches, not as extremes of a continuum, but as two different modes used deliberately for different purposes. Consequently, an artificial perception model that distinguishes between two different input modes is needed.

To automatically replicate human perception, it is essential to understand how humans interpret engineering sketches. Scientific strategies towards this goal typically rely on experimentation and interviews with subjects in order to make perceptions explicit (a recent example can be found in [Pac05]). However, understanding perception

is different from replicating perception, as the latter requires an “algorithmic” knowledge of the process. This paper describes the guidelines and hypothesis of a new algorithmic model for the artificial perception of junctions that depict vertices and discusses a set of experimental studies designed to inform its development.

2. Terminology

The input for Sketch-Based Modeling (SBM) approaches are *sketches*, which are freehand drawings comprised of strokes. *Strokes* are scribbled lines sampled by a set of consecutive nodes caught between “pen-down” and “pen-up” movements. The result is an ordered sequence of points, which are connected by segments to approximate the original scribble. A vectorized line can be obtained by suitable fitting-lines-into-stroke approaches. An example of a sketched line is shown in Figure 1.




Fig. 1. Sketched lines are stored as strokes and parsed to produce vectorized lines.

If the line-drawing depicts a polyhedral shape, then lines represent edges, with some circuits of lines representing faces. Vertices, which are the points of the shape where edges meet, are represented by junctions, which are the points of the drawing where line tips meet.

The input for our vertex-merging approach is a set of vectorized line segments (or simply “lines,” as referred to henceforth) delimited by their tips. Vertex detection must merge dangling tips to determine junctions that depict valid vertices (Figure 2).



Fig. 2. Merging of three tips (left) to form a single junction (right)

The output is a *line drawing*: a list of lines and a list of junctions, where each line connects two junctions. Junctions are x, y coordinate pairs of shared endpoints that likely correspond to vertices of the depicted object.

We distinguish *careful* and *casual* sketches (also called *detailed* and *quick* sketches). Examples of both types of sketches are provided in Figure 3. Based on the classification of sketches defined by Ferguson [Fer92], we assume the idea that thought sketches are usually casual, while prescriptive sketches tend to be more careful.




Fig. 3. Careful (left) and casual (right) sketches of a prism

3. Related work

Although some vectorization applications are available (e.g., CALI [JF99]), they are mostly valid for accurate and refined sketches. Our goal is to merge tips of vectorized lines to produce 2D junctions that depict 3D vertices. To this end, our approach ignores common methods for analyzing digital images which capture real-world scenes, as those images are *rich* in content (i.e. they contain a large amount of information) but store information at a *low-semantic level* (e.g. raster or bitmap images). In contrast, SBM inputs are *sparse* images (which contain a reduced

set of strokes) that can be converted to *high-semantic* line drawings (containing lines and junctions) by applying perceptual principles. Furthermore, most methods for analyzing 2D camera images only consider small errors (such as those resulting from lens imperfections and electronic noise), while sketching errors are typically more significant, as “imprecision allows designers to work with rough values and avoid premature commitment” [JGH09].

Approaches for drawing vectorization and beautification typically include *endpoint snapping* strategies (like the intersection detection in [GFP13]). A recent contribution by Fišer et al. [FAS16] goes one step further, as it builds on the idea of chaining different beautification rules. The authors argue that the rules do not depend on each other, while their combined effect can lead to complex modifications of the input stroke. The idea looks promising, despite the debatable independence between rules. In fact, evidence on the mutual relationship between reconstruction rules is well described ([LF92], [LS96], [CCC04], [YTJ08]). Furthermore, even the sequence selected to chain the rules together may affect the outcome. For instance, two lines that are considered within range for being nearly perpendicular may get out of range after one of them is slightly rotated to snap its endpoint.

In this context, it can be argued that optimization based formulations [FLB16], or machine learning techniques (such as convolutional network approaches [SIS16]) are better strategies. Nevertheless, these approaches necessarily rely on global perception. Training a machine learning system with abstract drawings that do not represent any perceivable shape would result in trivial training, while using sketches that depict shapes necessarily forces the trainer to use—even unconsciously—her global perception capability to judge the particular junctions. This prevents emerging the *local* perceptual rules that also govern the human detection of vertices that we want to replicate. In other words, we advocate for an algorithmic approach that detects junctions at a local level, considering just the neighborhood of the junction, and ignoring (as much as possible) the global perception of the depicted shape. This paper builds on the idea of formulating better rules for isolated problems before trying to combine them. Our goal is to improve the approach of determining a linearly decreasing likelihood for tip snapping based on the distance between tips [FAS16].

To this end, we build on the perceptual rules described in the literature. A similar approach was recently proposed by Liu et al. to simplify over-traced sketches [LWH15]. The chicken-or-the-egg dilemma described by Liu et al. can be reformulated as follows: vertices are defined by stroke endpoints, and stroke endpoints must meet at vertices. Therefore, similarly to Liu et al., we adopted a strategy based on not only proximity and continuity, but also closure.




Fig. 4. Junctions detected by proximity (left) and closure (right)

Detecting whether a user is drawing carefully or casually is feasible—see, for instance, recent efforts to use the design intent embedded in the features of the product to quantify its uncertainty [ECE16]—but sketches may contain a mixture of casual and careful junctions. Therefore, we assume that both types of junctions must be detected and adopt an approach that uses two stages. In the first stage, adjacent tips of lines are reified into common junctions, while in the second stage intersecting lines are trimmed (or extended) to detect emergent junctions.

Regarding perceptual principles and theories, the goal of this paper is not to develop new theories or criticize existing ones, but to replicate human perception algorithmically. To this end, we use perceptual principles and rules that can describe human perception reasonably well (e.g., Hoffman rules [Hof98]) and are capable of producing algorithms.

4. Human perception of sketched junctions

Our general vision is that, for complex shapes, it is easier for computers to perceive at a local level first, and then recombine the more likely local perceptions at a global level. This is because local detections reduce the workload of the subsequent global stage. The more trivially detected local perceptions become the less complex the global perception will be. Yet, as it seems to occur with humans, the procedure is likely to be iterative. If the previous experience of a subject helps her perceive a particular shape (e.g., a cube), then finding its vertices is straightforward. However, when the same subject is confronted with an unknown shape, determining its vertices comes first, as it greatly reduces the complexity of the subsequent perception of the global shape. If only some of the junctions are perceived clearly, then the subject can attempt to make sense of the global shape, to later check if her global perception endorses the previously detected junctions and helps her fix the remaining.

According to this vision, we intend to score the likelihood of the junctions in a sketch. We hypothesize that sketches convey perceived geometry instead of actual geometry. Therefore, we do not assert “this is a vertex” but “this is more or less likely to be a vertex.” By doing this, we provide information for a future global perception mechanism that may take advantage of the local interpretation approach described here. This future improvement is based on our hypothesis that humans put models in the context of scenes, and those scenes influence the perception of shape. Therefore, a local intersection of a set of strokes which can barely be accepted as depicting a common junction when perceived separately, may be seen as a valid vertex when put into context (see, for instance, the lower right junction in the casual sketch of Figure 3).

Our local perception approach is rooted in two complementary principles. Following the Gestalt law of proximity, two or more carefully drawn tips that are close to each other are to be reified as a single junction (Figure 4, left). The way this rule applies is explained in detail by Hoffman’s rules #2, “if the tips of two lines coincide in an image, then always interpret them as coinciding in 3D,” and #4 “Interpret elements nearby in an image as nearby in 3D.”

Similarly, if a sketch is assumed to be casual, tips are presumed to be imprecise, and the junctions must emerge, not at the tips of the lines but at the intersections between the lines (or their extensions). In this case, the law of closure dominates (Figure 4, right). As Hoffman’s rule #13 states, “if three or more curves intersect at a common point in an image interpret them as intersecting at a common point in space.”

However, these driving ideas are not enough to convert perceptual rules into an algorithm. A valid sequence needs to be found to apply the different rules and quantify the qualitative limits described by psychologists. To do so, we conducted a group of experiments to analyze how subjects perceive junctions in both carefully and casually sketched technical drawings. Junctions in careful sketches are uncontroversial—everyone (and most algorithms) will agree where they ought to be—but there is less agreement on junctions in casual sketches. The goal of our experiments is to determine the thresholds between careful sketches (everyone agrees) and casual sketches (most people agree), and between casual sketches (most people agree) and poorly drawn sketches (no general agreement). We note that the ultimate goal of these experiments is not the discovery of new perceptual principles or rules, but to establish a set of criteria and a benchmark to develop and validate our algorithmic approach.

5. Experimental method and results

Three experiments were designed to analyze how subjects perceive junctions in both carefully and casually sketched technical drawings.

Because our work focuses on engineering sketches, all subjects came from technical backgrounds. Our participant pool ranged from first year students of various engineering disciplines to experienced mechanical engineering professors. We hypothesized that there is no difference in the perception of vertices depicted by junctions in casual sketches between groups of subjects with different levels of exposure to technical drawings. This hypothesis was tested in Experiment #1.

In parallel, Experiment #0 validated a set of examples for benchmark use in future algorithmic developments. After analyzing Experiment #1 data, we hypothesized that people are more tolerant of imprecise vertices if these are perceived as part of a recognizable figure. Experiment #2 was designed to test this hypothesis.

5.1. Experiment #0

The goal of Experiment #0 is to demonstrate that there are no significant differences between our test sketches and those collected from other subjects.

To analyze what type of errors and inaccuracies people usually make when drawing casual sketches, we interviewed 91 subjects (17 from Spain, 61 from the US, and 13 from Italy). Most subjects had an engineering

background (18 were engineering educators and 73 engineering students). Every participant was asked to casually sketch the four models shown in Figure 5. These models are a subset of the models used in Experiment #1.




Fig. 5. Semi-transparent view of the models used in the experiment

To illustrate the task to be completed, the example shown in Figure 6 was provided to participants.




Fig. 6: Sample shape shown to participants as part of the instructions.

Participants were asked to draw each line in the sketch as a single stroke and include hidden edges by representing them as solid lines (not dashed).

Initially, all sketching was going to be completed digitally in a tablet device. However, we quickly realized that participants who did not have previous digital sketching experience seemed to be more concerned about the technology than the sketch itself (as described, for instance, in [OK11]) which clearly affected the results. Participants were then asked to create their sketches on paper, and then scan and email them to the research team.

The criteria used to parse or reject the sketches can be described as follows: (a) only fully sketched drawings were considered (sketches with missing lines were rejected), (b) dashed lines were treated as solid lines going from the tip of the first dash to the tip of the last dash, (c) short beautification strokes added by some participants to better fit the end of imprecise strokes were ignored, and (d) sketches with overtraced strokes were also rejected. Examples of all four types of incorrect sketches are shown in Figure 7.

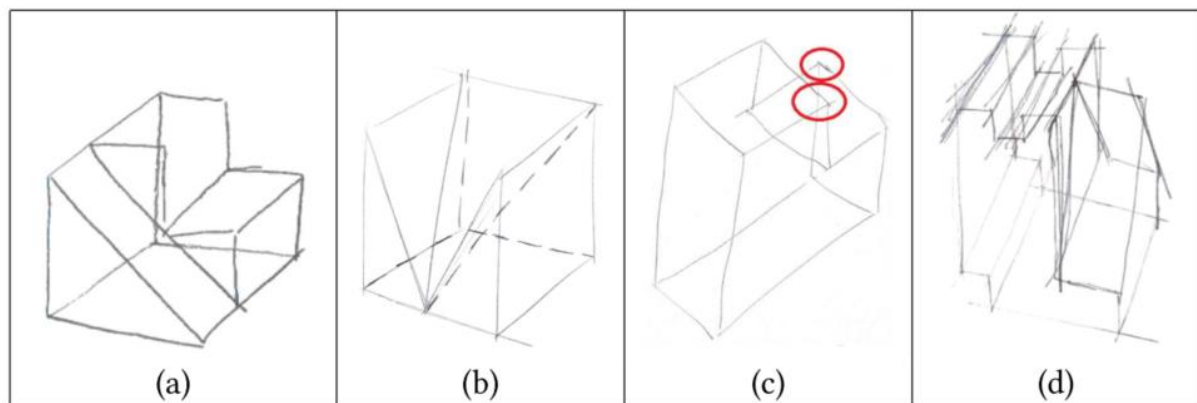


Fig. 7: Examples of incorrect sketches (a) missing strokes, (b) dashed hidden lines, (c) ignored short beautification strokes and (d) overtraced sketches.

A total of 44 sketches were rejected for model 1, 25 for model 2, 24 for model 3 and 43 for model 4. All rejected sketches came from engineering students. We speculate that sketches of Model 4 frequently included missing lines because of the complexity of the shape, while missing lines in sketches of Model 1 were possibly due to the model's point of view or orientation.

From our analysis, we conclude that the drawings used in the perception experiment are representative, as their coverage of junction types does not differ significantly from the sample of collected drawings. We are aware of a possible cultural bias in subtle aspects of the sketching style, which would require further examination.

The examples used in this experiment can be useful to initially validate an algorithm aimed at measuring the merit of a group of tips to be perceived as defining a common junction. We define a figure of merit in the range [0, 1], where 1 means that the junction is to always be perceived as such by humans, whereas 0 means that humans will not perceive the junction.

After analyzing the valid responses (228 sketches from 4 different models), the following observations were noted, which can be used to guide the design of an algorithmic approach to vertex perception:

1. Common errors and inaccuracies related to the junctions drawn were identified:
 - a. Existence of gaps between tips of lines that should intersect with each other. Gaps were detected in 138 of the 228 sketches (60.5%). Examples are shown in Figure 8, column a.
 - b. Extended lines that intersect with each other. This inaccuracy was detected in 149 sketches (65.4%) (Figure 8, column b)
 - c. Overlapping between junctions that could belong to different vertices of the model. This characteristic was found in multiple sketches of models 1 and 4, perhaps due to the particular point of view of model 1 and the complexity of model 4. A total of 21 of 47 sketches of model 1 (44.7%) and 13 of 48 sketches of model 4 (27.1%) had lines of different vertices which seem to meet at the same junction (Figure 8, column c).
2. Despite explicitly requesting casual sketches, some participants created sketches that contained a mixture of carefully and casually drawn junctions.
3. We assumed that there were no significant sketching differences between left and right-handed participants to justify an additional experiment (13 of the 91 participants were left-handed). Likewise for the participants' national origin and cultural background. Additionally, although we acknowledge a lack of diversity in our sample, we consider sketching differences between people that tend to draw lines from left to right and vice versa negligible, as these tendencies would only imply an increase of casual junctions in the right or left sides of the sketch.

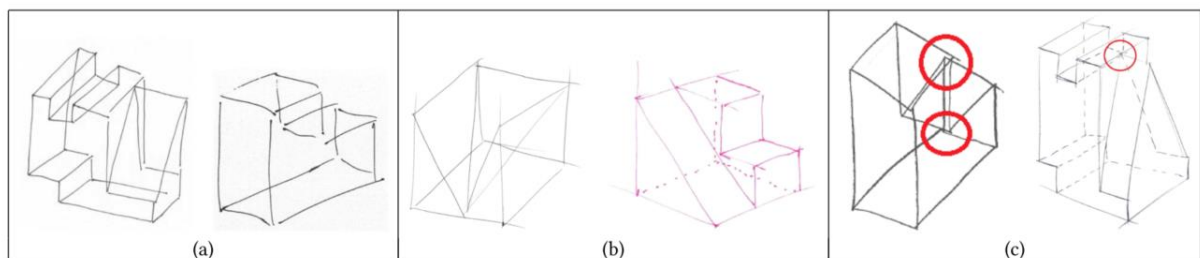


Fig. 8: Common inaccuracies in participants' sketches, (a) gaps between line tips, (b) extended lines (c) overlapping junctions.

5.2. Experiment #1

For this experiment, we hypothesized that subjects generally perceive the same junctions in casual sketches, regardless of their experience and expertise in technical drawing. A group of 12 casual sketches were created by the authors and distributed to two groups of participants (students and instructors). The geometric shapes were deliberately selected to include "relevant" features (Figure 9). For each sketch, participants were asked to number the perceived junctions and highlight the tips of the lines that intersected at that junction. All perceived junctions are labeled in Figure 9 as "sketch number" followed by a "V" plus the "junction number."

38 questionnaires were collected from undergraduate students and 38 questionnaires from instructors of engineering graphics, mechanical engineering and architectural graphics.

By comparing the relative frequency with which students and teachers perceived each joint as a vertex, two distinct groups of sketches emerged. In the first group (sketches 1, 9, 10 and 12), all vertices in the drawings were perceived successfully by more than 90% of the subjects. Similarly, in the rest of the sketches most vertices were also perceived with a high percentage of consensus by both groups. Only a few joints were perceived poorly (by both groups of participants). This poor recognition can be due to different factors: an ambiguous point of view of

the sketch (3V7, 7V9, 8V7); the existence of gaps between the tips of two lines (4V7, 4V8); the closeness between two tips of different lines (6V13); the intersection of two long lines (4V9, 4V10, 6V14, 6V15, 7V7, 7V8); or the complexity of the sketch (5V16, 11V24).




Fig. 9. Casual sketches and perceived junctions in experiment #1

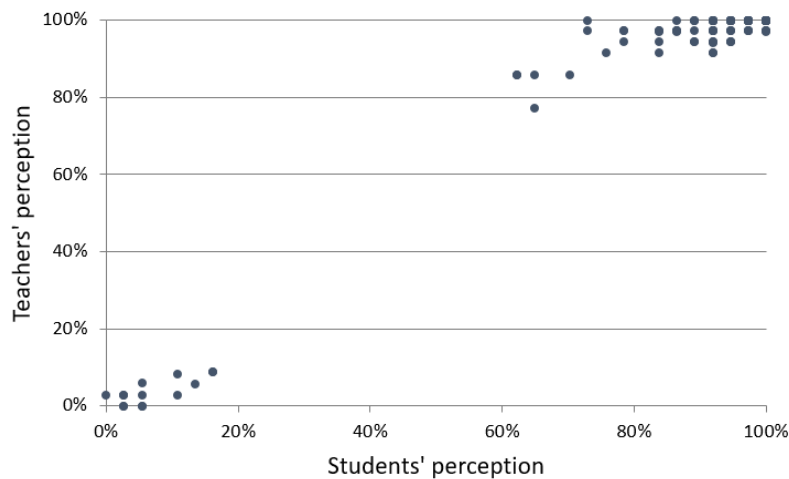


Fig. 10. Perception of vertices (Teachers vs. students) for the 12 sketches used in our experiment

The relationship between the students' and teachers' perception of all the 140 vertices (distributed across the 12 sketches shown in Figure 9) is illustrated in Figure 10. The scatter plot shows consensus in perception between the two groups, both for easily perceived vertices as well as those that are poor recognized.

Results of the Pearson correlation corroborate a significant positive association between the degree of perception of students and teachers for all 140 vertices, $r(140) = .977, p < .01$. Therefore, it is reasonable to combine the responses from all subjects to further analyze their behavior when detecting junctions.

Despite some consensus, not all vertices are perceived equally. By observing the junctions in Figure 9 (which were perceived with high agreement levels), we hypothesize that there is consensus in the perception of junctions

that represent trihedral vertices, and that people agree to perceive a junction that clearly represents a “squared” vertex (where two orthogonal lines meet).

To this end, we analyzed the data obtained from the experiment when subjects marked the tips of the lines that intersected at the junctions they perceived. The frequency with which subjects perceived a junction and the specific number of strokes intersecting at that junction is illustrated in Figure 11.




Fig. 11. Frequency of junction perception based on the number of lines marked by subjects

These results seem to validate our hypothesis: (1) subjects’ perception is not random and (2) people apply Hoffman’s rule #13 and perceive trihedral vertices when observing a junction shared by 3 strokes. Although a conclusive validation would require an additional experiment, we decided to include these observations as a design criterion for our algorithm and indirectly validate whether it replicates human perception.

Although junction perception seems to decrease when subjects recognize two lines sharing the junction, a special case can be identified: figures representing flat shapes with right angles. When the sketch represents a clearly flat figure (such as Example 2 in figure 9), the frequency of perceiving squared vertices increases to 75.6%. Otherwise, the gaps between the tips of the strokes seem to discourage subjects from marking more than two lines (such as 6V4, 6V7, and 6V10). Only when the sketch clearly represents a flat figure can we conclude that people agree to recognize squared vertices (Example 2).

We also observed that junctions with high dispersion in responses are related to junctions with more than three intersecting strokes, where some of the converging lines could be collinear (3V6, 5V24, 9V4, 9V5, 9V9, 10V9 and 11V12). This behavior reinforces our previous assumption that collinearity prevents people from perceiving junctions.

5.3. Experiment #2

The objective of this experiment was to determine the thresholds between careful sketches (everyone agrees) and casual sketches (most people agree), and between casual sketches (most people agree) and poorly drawn sketches (no general agreement).

We consider that the level of quality of the sketches also influences how people perceive the junctions. In this sense, we hypothesized that people are more tolerant of imprecise vertices if these are perceived as part of a recognizable figure.

To validate this hypothesis we used 12 new sketches by creating variations of examples 2, 10, and 12 from Experiment #1. Participants were asked whether they perceived a flat figure or a particular polyhedral 3D model. They were also asked to number the vertices they perceived and mark the lines intersecting at those junctions.

A total of 50 responses were collected for each example. The 12 examples and the relative frequencies with which each sketch was recognized as a flat figure or a 2D representation of a polyhedral 3D model are shown in Figure 12. The relative frequency with which each junction was detected in the drawing is also shown.

Do you perceive a square?			
Yes (88%)	Yes (59%)	Yes (50%)	Yes (41%)
Do you perceive a square pyramid?			
Yes (63%)	Yes (84%)	Yes (63%)	Yes (50%)
Do you perceive a cube with an inverted pyramidal hole?			
Yes (72%)	Yes (78%)	Yes (53%)	Yes (34%)

Fig. 12. Casual sketches and perceived vertices in experiment #2

Results from Pearson correlation indicate a significant positive association between the subjects' perception of 3D models depicted in the drawings and the frequency of junction detection, $r(70) = .77, p < .01$ (Figure 13).

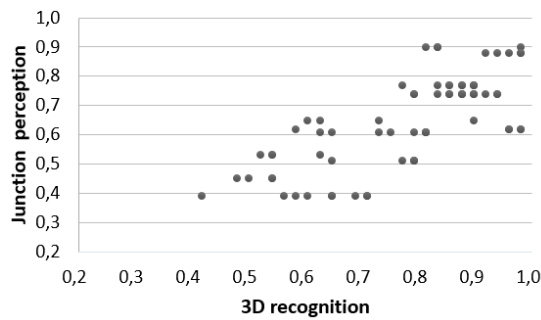


Fig. 13. Scatter plot of 3D model recognition vs. junction detection

These results are consistent with our hypothesis that the more clearly the flat figure or the polyhedral 3D model is perceived in the drawing, the more consensus there is in detecting junctions. People are more tolerant to imprecise junctions when they apply global perception. Therefore, a local perception approach should be less tolerant than a global one (even at the risk of ignoring some actual junctions). Finally, Experiment #2 allows us to determine the extent to which inaccuracies at intersections between lines affect the detection of junctions.

We tabulated the relative distance of each junction, calculated as the ratio between the distance from the intersection point (or from the centroid of several intersection points) to the most distant tip; and the maximum length of the lines sharing that junction.

It was observed that for relative distances greater than 11%, junction detection for flat figures decreased significantly and the responses regarding the number of lines sharing the junction were greatly dispersed. Similarly, when intersections occurred at line segments, junction detection also decreased or introduced great dispersion for relative distances greater than 25%. In the middle of the interval, the analysis was not so clear, as people seemed to have used both global and local perception.

We conclude that the detection of vertices in engineering sketches seems to be related to the application of global perception strategies when interpreting the sketch. However, the parameter “Distance” clearly influences this perception even if only local perception applies. Thus, this parameter—which allows for a certain degree of flexibility when detecting junctions—was taken into consideration in the design of our algorithm, which is required to find junctions without prior knowledge of whether the sketch is careful or casual, or what object is being represented (even if the scope is currently limited to polyhedral shapes).

6. Additional analyses

Additional analyses were conducted to determine how people perceived junctions in Experiment #1. These observations helped guide the design of an algorithm whose effectiveness had to be checked after the fact.

In Table 1, the column “Qty junctions” represents the quantity of junctions that were perceived most frequently in each example of Figure 9; the “% observed” column represents the relative frequency with which people marked this quantity of junctions in each example; the third and fourth columns show the frequency with which subjects marked fewer or more junctions respectively.

Table 1: Summary of junction detection per example

Example	Qty junctions	% observed	% less junct.	% more junct.
1	8	93.6%	6.4%	0%
2	4	75.6%	24.4%	0%
3	6	78.2%	12.8%	9%
4	6	84.6%	3.8%	11.5%
5	26	62.8%	36.2%	1%
6	12	89.7%	6.4%	3.8%
7	6	72.4%	10.5%	17.1%
8	6	96.1%	2.6%	1.3%
9	14	84.6%	15.4%	0%
10	9	90.9%	9.1%	0%
11	24	64.1%	32.1%	3.8%
12	5	93.5%	6.5%	0%

We can conclude that:

- People tend to detect junctions more easily when lines clearly intersect as opposed to when their extensions intersect. For instance, in Example 2, 63.2% of the participants who perceived fewer than four junctions, perceived them mainly where direct line intersections occurred (2V1) but not at intersections of line extensions (2V2-2V4). Similarly, although 93.5% of participants detected five junctions in example 12, no consensus was reached for 12V5.
- People rarely perceive junctions between long lines that cross each other. For instance, no subjects marked a junction at the intersection of 1V2-1V6 and 1V1-1V4, whereas only 3% of the subjects in Example 4 detected junctions at intersections between long lines (4V9 and 4V10). In example 6, not a single subject marked a junction between intersecting lines defined by 6V2-6V5 and 6V3-6V4, and 6V14 was marked only by 1% of subjects.
- Collinearity prevents people from perceiving junctions, as it occurred in 3V6 which is shared by several (some almost collinear) overlapping strokes and has the lowest perception percentage of junctions in example 3.
- People tend to look for trihedral junctions when overlapping occurs and the interpretation of the drawing as a 3D model is unclear. In example 3, 9% of participants duplicated junction 3V6 (3V7) by detecting seven junctions. In these cases, subjects seemed to apply Hoffman’s rule #13 [Hof98] and prioritize the search of trihedral junctions. The same occurred in example 8, where 1% of subjects detected seven junctions because 8V6 junction was duplicated (8V7). In example 10, the most under-detected junction was 10V9.
- Junction perception seems to depend on how clearly the depicted 3D model is interpreted. For instance, in Examples 1, 8, and 12, most subjects detect the same quantity of junctions in the drawings. Nevertheless, in examples 2, 3, and 7 the agreement decreases to 75.6%, 78.2% and 72.4% respectively. In these cases, gaps between tips, overlapping strokes added to particular points of view like examples 3 and 7, seem to increase the misunderstanding of the drawings. In the case of 7V7 and 7V8, the

perception of junctions between long intersecting lines reached 12%. Once again, the point of view of the drawing did not facilitate the interpretation of the depicted object. Therefore, when the interpretation of the 3D model is unclear, people seem to apply local perception instead of global, to detect junctions that may represent 3D vertices.

- Finally, complex examples with many junctions, such as examples 5 and 11, produced the lowest levels of agreement. Only 62.8% and 64.1% of participants detected 26 and 24 junctions respectively. We speculate that the large number of lines and junctions distracted subjects and significantly impaired perception.

7. Algorithm

Two main ideas are combined in our algorithm: (1) well-known Gestalt Principles must guide the *perception* of the junctions, and (2) the strategy must manage both *careful* and *casual* sketches [CVP19].

Careful sketches are based on reification: the action of perceiving the tips of two or more lines as representing one single junction. By applying the Gestalt law of proximity, two or more carefully drawn tips that are close to each other are to be reified as a single junction. The way this rule applies is explained in detail by Hoffman's rule #2, "if the tips of two lines coincide in an image, then always interpret them as coinciding in 3D", and #4 "Interpret elements nearby in an image as nearby in 3D" [Hof98].

For casual sketches, the driving idea is emergence, i.e., tips are presumed to be imprecise, thus junctions must emerge, not at the tips of the lines but at the intersections between the lines (or their extensions). In this case, the law of closure dominates. As stated by Hoffman's rule #13, "if three or more curves intersect at a common point in an image interpret them as intersecting at a common point in space" [Hof98].

The goodness of the detected junctions must be quantified by a figure of merit, which is in the range [0, 1]. 1 means that the junction is to always be perceived as such by humans; 0 means that humans will hardly ever perceive the junction.

The full source code of the algorithm is freely available at [CVP17], including examples that demonstrate the capabilities and limitations of the approach.

7.1. Reification of careful vertices

The classic nearest-neighbor criterion of merging tips by determining their distances to one another is used by our approach to merge junctions that represent vertices defined by carefully drawn lines. The method was improved in a number of ways.

First, we consider that longer lines usually belong to main features—which are perceived first—whereas shorter lines are usually perceived as belonging to secondary details. Thus, the algorithm searches for junctions between longer lines first.

Second, errors that are acceptable in the location of the tips tend to be proportional to the total length of each line: errors in the location of the tips of longer lines are accepted to be greater in absolute value than those of shorter lines. Thus, the threshold distance between the tip and the candidate junction is defined as a percentage of the length of each line.

Third, since the sketches are assumed to depict orthographic representations of flat figures or pictorial representations of polyhedral shapes, the range between right angles and cubic corners (see [Per71] for a definition of cubic corners) is prioritized. To this end, an anisotropic *merging area* is defined around each tip (see Figure 14 left, where the tip v_2 of line e_2 must be closer than v_3 to be labeled as close enough to merge with v_1 .) The merging area depicted in Figure 14, left is merely illustrative, since the intersection between the lines may not necessarily be coincident with any of the tips. Therefore, the distance allowed between two lines (e_1 and e_2 in Figure 14, right) is calculated and then compared to the actual distance between their tips. The allowed distance (r_2 in Figure 14, right) is calculated as the product of a threshold *maxDist* (fixed at 8% of the length of the line) and an *allowance* = $2 - \cos(e_i, e_j)$. The allowed distance ($maxDist * allowance$) is at its maximum value for 90 degrees (20%), and still high for 60 degrees (15%). In axonometric views, squared lines are commonly depicted by lines at angles close to 60 degrees.

The allowance penalizes merging tips of collinear lines that are poorly drawn. Naturally, collinear lines are uncommon in polyhedral vertices and non-collinear lines are seldom depicted as collinear (as this situation would conflict with Hoffman's rule #3 "Always interpret lines collinear in an image as collinear in 3D" [Hof98]). By using the allowance to calculate the merits of the candidate junctions, the merging of perpendicular lines is prioritized.




Fig. 14. The neighborhood of each tip depends on the relative angles of the lines that contain the candidate tips to merge (left), while close tips are merged at the intersection of both segments, even if the intersection falls outside the merging area (right)

If tip v_2 is sufficiently close to v_1 (this is, if $|v_1 - v_2| < |r_2|$, where $|r_2|$ depends on the relative orientation between e_1 and e_2) the intersection of the two segments is defined as the common merged tip, regardless of whether or not this new tip falls outside the merging area (Figure 14, right). This way, lines are prevented from rotating while merging their tips.

Up to this point, our approach assumes that tips of vectorized lines are not yet connected to each other, i.e., the *valence* (number of lines that share the tip) is 1 for all the tips. However, this is not the case when the merging-vertices algorithm progresses, as tips that have already been visited may be partially merged. Therefore, the goal is to prioritize the most populated junctions.

The general rules to *cluster* tips of valence other than 1 are as follows: (1) If valences are different for both tips, the tip with the highest valence remains fixed while the other is moved to the cluster, (2) If valences are equal (and different from 1) both tips are moved to the cluster into the midpoint they define. A clustering threshold is defined to prevent lines from excessive rotation. A merge is declared valid if all the lines connected to the moving tip rotate less than $maxRot$ (fixed at 10 degrees) when redirected to the new tip.

Therefore, the merit of reified junctions is assigned -1 if tips are separated more than the allowed distance or the maximum rotation of any line exceeds $maxRot$. Otherwise, the merit of reified junctions is in the range average-to-good, calculated as a fixed minimum reward for careful junctions RM (fixed at 0.50), plus two variable components. The first variable component, which depends on the distance merit, equals RD if both tips are coincident, and decreases linearly to 0 for tips separated by the allowed distance ($maxDist * allowance$). The second variable component depends on the rotations: it equals RR if no line rotates to merge the tips, and linearly decreases to 0 if the maximum rotation of any line equals $maxRot$. To ensure that the merits of valid junctions are in the range $[0, 1]$, RD is fixed as $(1 - RM)$, and RR is null, for lines with valence 1. For other lines, both RD and RR are fixed as $(1 - RM)/2$.

7.2. Emergence of casual vertices

Two rules apply to help the emergence of junctions at the intersections between casual lines (or their extensions).

First, Hoffman's rule #13 is used to prioritize trihedral junctions as follows: three dangling lines that intersect close to each other, and close to their tips, define a new junction. The maximum allowed distance between intersections is $maxDistTriplets$ (set to 10% of the length of the longest line in the triplet). The distance from the intersection to the tips must fit in a *valid range*. It may not be longer than $inTol$ (50%) of the average length of the three candidate lines for intersections inside the segment, and no longer than $outTol$ (25%) of the average length of the three candidate lines for intersections in the extension of the segment. Assigning an asymmetric range ($inTol > outTol$) prioritizes intersections inside the segments, which was experimentally determined to be the common perceptual behavior.

The intersections must be the closest to the tips in order to prevent false merges such as the one illustrated in Figure 15, left. In this case, semi-line e_1 could be incorrectly merged with semi-lines e_0 and e_2 if they were processed from longer to shorter. The three semi-lines that intersect close to each other are connected to a common junction located at the centroid of the three intersection points (Figure 15, middle). The merit of the emerged triple junction is in the range $[0, ETM]$, since it is assigned as $ETM * (1 - distCentroid / maxDistTriplets)$, where $distCentroid$ is the distance between the centroid and the most distant intersection point. In our implementation, ETM is fixed as 0.5, to ensure that the merits of emerged junctions are in the average-to-poor range. Next, a relaxed search is performed after removing the condition that intersections must be the closest ones to the tips. This provides a second chance

for detecting those casual triplets that are close to each other and include long lines (which are parsed first) that are more casual than the short ones.




Fig. 15. Triplets should only merge semi-lines that intersect closer to the tips (left), while the centroid of the intersections between three dangling lines emerges as a junction (middle), and a line that needs excessive rotation to merge to the original neighboring tips can merge when relaxing rotation (right)

Second, emergence is also considered by clustering dangling lines to nearby junctions, i.e., a dangling line that passes closely to a junction of valence two or higher must be clustered to it. “Passing closely” implies that the distance between the line and the junction (such as d_3 in Figure 15, right) is less than $maxDistDangling$, which is set to a percentage (10%) of the length of the dangling line. An allowed distance ($maxDistDangling * allowance$) is used, where the *allowance* is calculated between the dangling line and its most collinear line in the junction (i.e. lowest allowance between the dangling line and each of the lines in the junction). It also implies that the point of perpendicular projection of the junction on the semi-line must fit in the *valid range*: $inTol$ of its length from the tip of the semi-line if it is inside the segment (t in Figure 15 right), or $outTol$ if it is in the extension of the segment.

The merit for dangling lines merged to junctions (*EDM*) is calculated as the minimum of two components: first, merit is 1 for dangling lines passing over the junction, and decreases down to 0 for lines passing at a distance of $maxDistDangling$. Second, merit is 1 for dangling lines where the closest point to the junction is the tip of the dangling line, and decreases linearly for lines whose closest point to the junction is at the limit of the *valid range*. Dangling lines with negative merit are not merged. As a result, the new merit assigned to the junction where a dangling line is merged is the minimum between the current merit of the junction and the merit of merging the dangling line to the junction.

8. Validation

To test the validity of our approach, we checked the examples used in experiment #1 and concluded that the algorithm accurately detects what humans perceive (Figure 16). The only junction that was not fully detected by the algorithm with the default configuration was junction 12 in example 6, where tip 4 was not merged.

Examples with a large number of lines were not included in the previous figures so the junctions could be clearly displayed.

In fact, the increase in the number of junctions is not a problem for the algorithm, although it was noted during the experiment that people tend to get tired and lose some junctions when faced with sketches containing a growing number of strokes. Thus, we cannot test the validity of the algorithm compared to a statistically validated human perception for sketches populated with increasing number of junctions. Still, when a reduced group of persons was asked to agree or disagree with the junctions detected by the algorithm, the answer was always positive (although sometimes accompanied by surprise, as the algorithm detected junctions unnoticed by some people at first sight).

To test the merits provided by the algorithm, we selected careful and casual sketches of 2D shapes (see Figure 17 for some representative examples), as well as careful and casual sketches of 3D polyhedral shapes (see Figure 18 for some representative examples).

In most of the examples we tested, the perceptually correct solution was successfully found by the algorithm. Generally, neither false junctions were detected, nor close but different junctions were merged. In the worst cases, only certain junctions were incompletely merged. Furthermore, the merits seem to be in accordance with the probability of the junctions to be perceived as such by humans. However, the algorithm is clearly sensitive to the thresholds defined as parameters. For this reason, some intersections perceived as junctions by humans are not

perceived correctly by the algorithm when tuned with the default parameters (which were defined to prevent false positives).

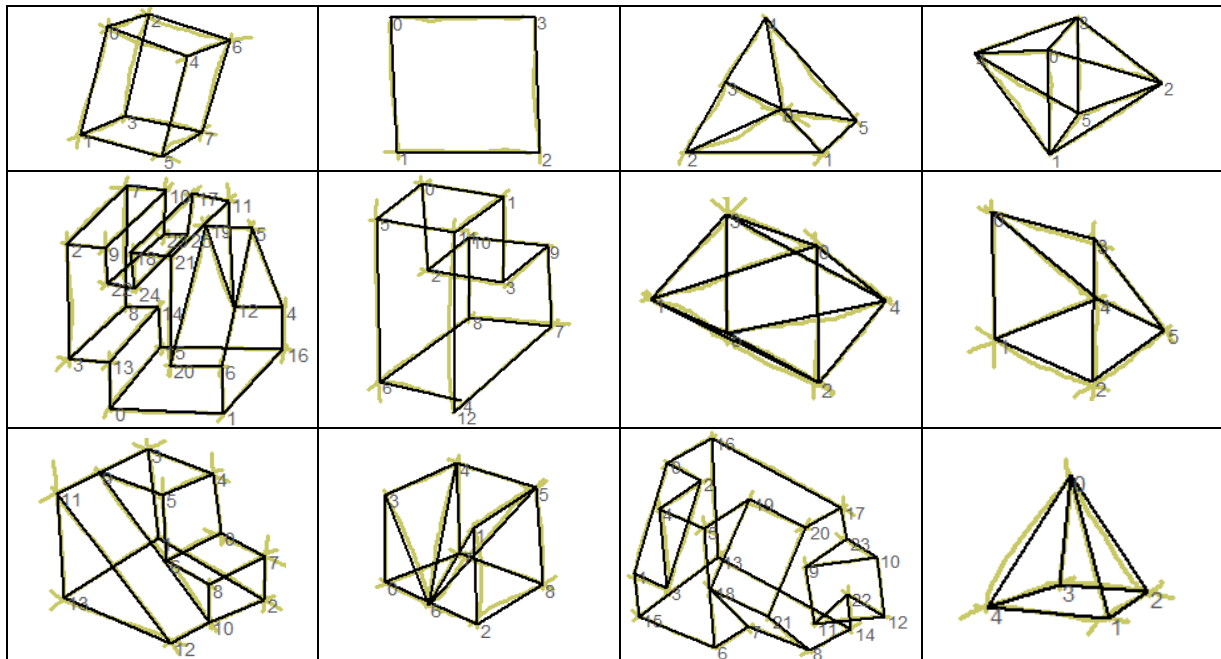


Fig. 16. Figures of the experiment, as detected by the algorithm

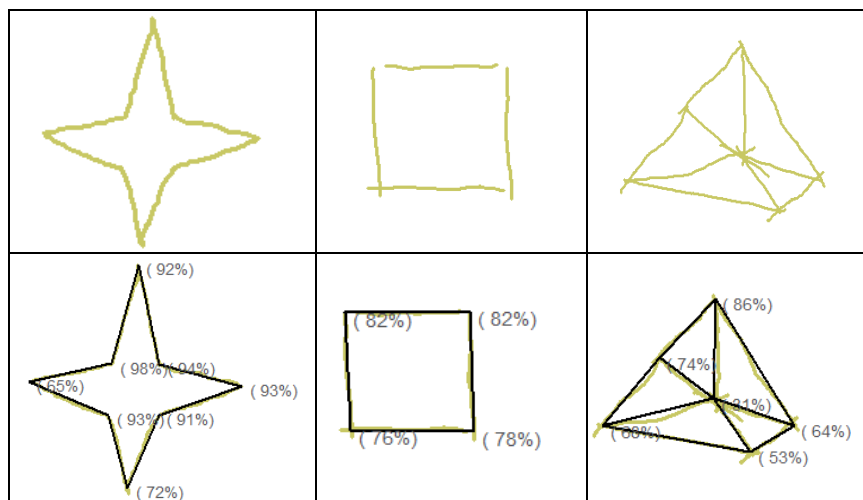


Fig. 17. Careful and casual sketches of simple 2D shapes, with their merits

Tips of exceedingly casual sketches can still be safely grouped—thus reducing false negatives—if the thresholds are increased, which is acceptable only if different corners are not too close to each other. For instance, by increasing parameter *maxDist* up to 13%, the casual star from Figure 19 can be solved correctly. The same occurs for example 6 in Figure 16, which is correctly solved by increasing *maxDist* to 10%. Similarly, example 3 in Figure 18 is correctly solved by increasing *maxDist* to 14%. However, arbitrarily increasing the parameters is prone to produce false positives, as shown in Figure 19, where the parameter *maxDist* was increased to 34% for the star, but barely 9% for the prism on the right side of Figure 3.

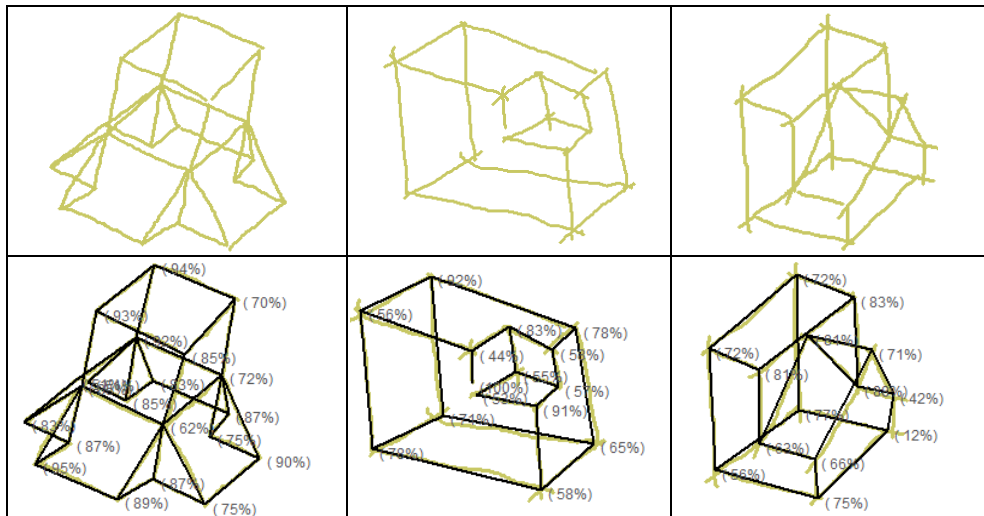


Fig. 18. Careful and casual sketches of 3D polyhedral shapes, with their merits




Fig. 19. Casual sketches solved with false positives

Reducing $maxDist$ (from 8% to 4%, for example) would prevent the algorithm from accepting casual junctions, while increasing it (up to 20%) would allow the algorithm to detect very poorly sketched junctions (but would also produce the incorrect merge of close junctions). Thus, the parameter may vary from less than 5% for careful and/or dense sketches up to 20% for those very casual but not dense (“dense” sketches are those with nearly overlapping junctions). Similarly, reducing the valid range by half ($inTol=25\%$ and $outTol=12\%$) prevents the algorithm from detecting false positives in dense careful sketches.

Varying threshold values is not recommended. Our belief is that most false positives (tips merged to define non-perceived junctions) will be hard to repair by users, whereas false negatives (junctions perceived by humans but not fully merged by the algorithm) can be clearly identified and are easy to edit manually (or in a subsequent automatic merging stage that could take advantage of global information). Still, while a balanced set of parameters is recommended, valid configurations for careful and casual sketches are tabulated in Table 2.

The first and second configurations are valid for the two types of sketches represented in Figure 3, as the first prioritizes reification, while the second prioritizes emergence (and both configurations output similar merits). Higher values of $maxDist$ are useful to detect lines that do not intersect close to each other (Figure 15 middle), but are counterproductive to detect lines with long tails (Figure 15 right), as the tips of the long tails may be close to an adjacent junction (Figure 19 right). We note that reducing $maxDist$ nearly cancels reification, thus preventing poorly drawn tips of neighbor junctions from emerging to produce false reified junctions.

We note $maxRot$ is independent on the size of the line, as we realized that humans perceive and are highly sensitive to even small rotations, irrespective of the length of the line. Furthermore, we observed that people tend to accurately determine the orientation of close lines, thus allowing the algorithm for large rotations to result in merging endpoints of lines intended to be parallel. Hence, in an attempt to avoid false positives, we opted for a strategy that may produce false negatives for short, isolated and poorly drawn lines.

Table 2: Recommended algorithm configurations

Parameter	Careful	Balanced	Casual
<i>maxDist</i>	12%	8%	4%
<i>maxRot</i>	5°	10°	10°
<i>RM</i>	0.2	0.5	0.8
<i>Valid range: inTol</i>	25%	50%	50%
<i>Valid range: outTol</i>	12.5%	25%	25%
<i>maxDistTriplets</i>	5%	10%	15%
<i>ETM</i>	0.2	0.5	0.8
<i>maxDistDangling</i>	5%	10%	15%

The established merits could be used to implement a dichotomous algorithm by simply considering valid those candidate junctions with merits greater than a fixed threshold (in accordance with Table 2). Thus, a threshold over 0.5 implies accepting only good junctions (which mainly include careful tips separated by less than $maxDist/2$), a threshold of 0.25 results in accepting good and average junctions (which mainly exclude casual tips separated by more than $maxDist/2$), and a threshold of 0 accepts all the merged junctions.

9. Conclusions

This work aims to contribute at converting sketches into line drawings by using perceptual principles that guide the search for junctions. Detection of vertices is an important and useful stage in Sketch-Based Modeling research. However, current methods for vertex detection are still ineffective when applied to casual sketches.

We described a series of experiments that were performed to gain “algorithmic” knowledge on how people perceive 3D vertices represented as 2D unions in casual sketches. From the results from Experiment #1, we conclude that subjects’ perception is not random and people perceive trihedral vertices when observing a junction shared by 3 strokes. Besides, we get a checked benchmark set to validate future algorithms. A small benchmark set useful to test tentative algorithms also resulted from Experiment #0.

From results of Experiment #2, we found that the parameter “Distance” clearly influences the. Thus, it was taken into consideration in the design of our algorithm.

Experimental data analysis helped guide the design of the algorithm to detect junctions in line drawings obtained by vectorizing casual sketches used by skilled engineers and designers to quickly convey rough ideas.

Our hypothesis that carefully sketched vertices must be reified, while casually sketched vertices must emerge, was indirectly validated, as it guided the design and implementation of an algorithm that improves on previous vectorization techniques, which were based exclusively on reification strategies.

The perceptual essence of the algorithm combined with sequential detection provides efficiency. The algorithm detects and fixes the best defined junctions, and only applies more complex calculations to poorly defined vertices.

Our algorithm is limited to line-drawings of polyhedral objects vectorized from sketches. But the procedure only uses information that is available in the neighborhood of the candidate junction. Besides, the algorithm calculates a figure of merit that estimates how likely a junction is to be perceived as such.

10. References

- [1] [Jen92] Jenkins D.L. The Automatic Interpretation of Two-Dimensional Freehand Sketches. PhD thesis. University of Wales, College of Cardiff, 1992.
- [2] [Pac05] Pache M.W. Sketching for Conceptual Design: Empirical Results and Future Tools. PhD thesis. Technischen Universität München (ISBN 3-89963-176-5).
- [3] [XCS14] Xu B., Chang W., Sheffer A., Bousseau A., McCrae J., Singh K. (2014) True2Form: 3D Curve Networks from 2D Sketches via Selective Regularization, ACM Transactions on Graphics (Proc. SIGGRAPH), Vol.33, No 4.
- [4] [Kan79] Kanizsa G. Organization in Vision: Essays on Gestalt Perception. Praeger, New York, 1979.
- [5] [JGH09] Johnson G., Gross M.D., Hong J., Do E.Y.L. (2009) Computational support for sketching in design: a review. Foundations and Trends in Human-Computer Interaction, 2 (1), pp. 1–93.

- [6] [Fer92] Ferguson E.S. *Engineering and the Mind's Eye*. The MIT Press, 1992.
- [7] [JF99] Jorge J.A. and Fonseca M.J. (1999) A Simple Approach to Recognize Geometric Shapes Interactively, GREC99. URL: <http://web.ist.utl.pt/mjf/publications/2004-1999/pdf/grec99.pdf>
- [8] [GFP13] Governi L., Furferi R., Palai M., Volpe Y. (2013) 3D geometry reconstruction from orthographic views: A method based on 3D image processing and data fitting. *Computers in Industry*, 64, pp. 1290–1300.
- [9] [FAS16] Fišer J., Asente P., Schiller S. and Sýkora D. (2016) Advanced drawing beautification with ShipShape. *Computers & Graphics*, 56, pp. 46-58.
- [10] [LF92] Leclerc Y, Fischler M. (1992) An optimization –based approach to the interpretation of single 449 line drawings as 3D wire frames. *International Journal of Computer Vision*, 9(2):113-136.
- [11] [LS96] Lipson H, Shpitalni M. (1996) Optimization-based reconstruction of a 3D object from a single freehand line drawing. *Computer Aided Design*, 28(8):651-663.
- [12] [CCC04] Company P., Contero M., Conesa J. and Piquer A. (2004) An optimisation-based reconstruction engine for 3D modelling by sketching. *Computers & Graphics*, 28 (6), pp. 955-979.
- [13] [YTJ08] Yuan S., Tsui L.Y. and Jie S. (2008). Regularity selection for effective 3D object reconstruction from a single line drawing. *Pattern Recognition Letters* 29 (10), pp. 1486-1495.
- [14] [FLB16] Favreau J.D., Lafarge F., Bousseau A. (2016) Fidelity vs. simplicity: a global approach to line drawing vectorization. *ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH 2016*, 35(4), No. 120.
- [15] [SIS16] Simo-Serra E., Iizuka S., Sasaki K., Ishikawa H. (2016) Learning to Simplify: Fully Convolutional Networks for Rough Sketch Cleanup. *ACM Transactions on Graphics (TOG) - Proceedings of ACM SIGGRAPH 2016*, 35(4), No. 120.
- [16] [LWH15] Liu X., Wong T.T., Heng, P.A. (2015) Closure-aware sketch simplification. *ACM Transactions on Graphics*, 34(6), 168.
- [17] [ECE16] Ekwaro-Osire S., Cruz-Lozano, R., Endeshaw H.B., Dias J.P. (2016) Uncertainty in Communication with a Sketch. *Journal of Integrated Design and Process Science*, 20(4), pp. 43-60.
- [18] [Hof98] Hoffman D.D. *Visual intelligence: how we create what we see*. W.W. Norton & Company, 1998.
- [19] [OK11] Orbay G, Kara L.B. (2011) Beautification of Design Sketches Using Trainable Stroke Clustering and Curve Fitting. *IEEE Transactions on Visualization and Computer Graphics*, 17(5), pp. 694-708.
- [20] [CVP19] Company P., Varley P.A.C., Plumed R., Camba J.D. (2019) Detection of vertices in sketched drawings of polyhedral shapes. LNCS 11401 (CIARP2018), pp. 376-383.
- [21] [CVP17] Company P., Varley P.A.C., Plumed R., Camba J.D. (2017) Source Code for Finding Vertices in 2D Line-drawings of Polyhedral Shapes. Technical Report. Regeo (2018), Geometric Reconstruction Group. Retrieved from <http://www.regeo.uji.es/FindingVertices.htm>
- [22] [Per71] Perkins D.N. (1972) Cubic Corners, Oblique Views of Pictures, the Perception of Line Drawings of Simple Space Forms. *Geometry and the Perception of Pictures: Three Studies*. Technical Report No. 5.