A Preliminary Assessment of the Role of Conceptual Salience in Automatic Sketching

Ahmed M. H. Abdel-Fattah, Wael Zakaria, Nashwa Abdelghaffar, Nohayr Abdelmoneim, Stefan Schneider, and Kai-Uwe Kühnberger

Computer Science Division, Faculty of Science, Ain Shams University, Cairo, Egypt; Institute of Cognitive Science, Osnabrück, Germany.
{ahabdelfattah,wael.zakaria,nashwa.abdelghaffar,nohayr}@sci.asu.edu.eg, {stefschn@uos.de,kkuehnbe}@uos.de

Abstract. We discuss and empirically assess the abilities of the normal person to both recognise conceptual salience of objects and generate free-hand sketches of these objects, investigating some underlying cognitive mechanisms that seem to be mainly responsible for these abilities. The ultimate goal is to employ human-in-the-loop results to implement a general-purpose automatic sketch recogniser that is guided by the way people operate on sketches to perform the same tasks. The aim of the article in hand is to contribute to answering two particular questions in this regard: does conceptual salience affect object recognition (when humans identify objects sketched by others)? and do specific parts of a sketch play more significant roles than others in generating a sketch of this object?

Keywords: Sketch recognition, concepts representation, spatial relations

1 Introduction

Automatic recognition of hand-drawn sketches by devices equipped with touch interfaces is foreseen to play an indispensable role in 'gesture-based' software applications. Enabling machines to automatically understand gestures, particularly various free-hand drawn strokes, and recognise (or better anticipate) the composing object, using only the most salient parts, would be a breakthrough. Think about the various scenarios in which we can utilise one or more of technological facilitators, such as Google© Glass®, Apple© Pencil®, Microsoft© Hololens®, Microsoft© Ink, etc. In the (not so far) future, we need technology to facilitate and immediately interpret what we draw or write, translating this into machine-understandable forms of sketch, and employing the interpretation for whatever assistance needed. Similar to what happens when alternative search suggestions are generated while one 'writes' an incomplete phrase, a sketch 'recogniser' may construct alternative sketch suggestions for an incompletely given sketch of an object being searched for. By coupling the readily available functionality with a gesture-based search engine, a sketch that one just started to draw can automatically (and quickly) be guessed (therefore, completed), offering alternative suggestions that dynamically differ as one continues to provide different parts of the sketch in the form of input strokes.

Automatic recognition or construction of sketches by machines becomes even more essential when we consider applications that are not only directed to assist people with specific needs or having special skills [Csapó et al., 2015], but also to help a normal person exert fewer efforts in performing mundane tasks quicker and more efficiently. Stylus pens are already commonplace –although its efficient usage is usually associated with graphic artists or design professionals. However, the potential of general usage appeals to people who still use pen and paper. Research shows that more than two-third of people still use pen and paper for at least one hour a day, and design recommendations from formative studies suggest that "natural" input modalities (e.g., voice and digital ink) could help to overcome the drawbacks of text entry on phones and PDAs (cf. [Andrew et al., 2009]). More recent surveys also show that people prefer pen and paper note taking for increased productivity and retention, with approximately 65 percent of a thousand of respondents attributing this to the simplicity of writing by hand (cf. `http://www.neolab.net/`). Further future applications can include support services in software systems for shortening the path through complex menus, automatic sketch generation for manuals and assembly instructions; a bridging approach between computer vision and conceptual reasoning; creative usage of sketches in e-learning contexts; or search services in large knowledge bases utilising input sketches.

1.1 Challenges: Sketch-based Computations in the Future

A normal person sketches and recognises free-hand drawn sketches easily and seemingly without complex reasoning, even if the sketches were not as precise as what an artist would draw, and even if the person does not have a high degree of language proficiency. Nevertheless, the acquisition of sketch generation and recognition abilities is not innate, but is rather an attentive process that needs to be learned. Moreover, when spatial relations are of a central concern, sketches become more suitable than language and allow to more easily draw on one's well developed spatial intuitions than verbal descriptions do. Recognition or retrieval of sketches by computational tools, on the other hand, is generally difficult, and requires long computations or simulation of complex mechanisms (e.g., spatial reasoning, matching, analogy making, abstraction, indexing, learning, etc.) that are not as intuitive as the humans' processing for sketch production or recognition. This is one of the major reasons why most (if not all) of the existing computational treatments (e.g., [Wang et al., 2011], [Hammond et al., 2010], and [Yuan et al., 2008]) of this kind of work is still limited to small-scale datasets (basically because a treatment is hard to match and inefficient to index different models). The automation is, nevertheless, indispensable, because gesture and touch interfaces are continuing to become mainstream human-computer interfaces (HCIs) that grow in usage day after day. In addition, with the advent of (and the recent advancements in) virtual and augmented reality, traditional input methods (e.g., keyboards and mice) are soon becoming legacy interfaces.

1.2 Challenges: Sketch-based Representation Requirements

What we refer to as a 'sketch' captures only conceptually relevant parts of an 'object', created or displayed by a normal person. Contrast this to what we usually refer to as a 'picture' or 'drawing' of the object, which is created by an artist and usually full of details, among other differences. A sketch also captures the spatial relations between the object's parts, making their treatment substantially different from classical image processing. Sketches thus can be considered as an intermediate level of abstraction between raw subsymbolic streams of sensory input on the one side and icons on the other [Abdel-Fattah et al., 2015]. We use the term 'conceptual salience' here to refer to the sub-concepts (in the representation) that correspond to salient features identifying the object (which itself is referred to as the main concept).

Using an automatic recogniser to computationally operate with sketches (employing recognition, enhancement, production, or completion), an artificial intelligence (AI) model should at least overcome certain knowledge representation and reasoning (KRR) issues. First, the model should employ efficient ways to modularly represent knowledge as conceptual entities (so that it can be worked with). To capture conceptually relevant parts of objects, the AI model needs to address the modular representation and manipulation of conceptual entities [Abdel-Fattah, 2014] that compose the main concepts corresponding to the objects of the sketches. Some researchers have even proposed image retrieval to be viewed as a knowledge representation problem, where structured objects are retrieved such that syntactic and semantic aspects play an important role [Sciascio et al., 2002]. Second, the model should be able to apply different reasoning forms. Deductive and inductive reasoning seem to be the primary needed forms. But spatial reasoning is needed as well to capture the 'spatial relations' between parts of the objects. The model needs also to be capable of recognising sketches drawn by any person, be they talented in drawing (i.e., an artistic expert) or not possessing special artistic capacities. The model should in general be based on psychological findings, because human recognition is not only data-driven, but crucially governed by cognitive mechanisms and principles, such as Gestalt principles [Dastani and Scha, 2003] and analogical reasoning [Gentner et al., 2001]. It seems, thus, next to impossible to develop AI computational models for simulating and utilising mechanisms of cognitive beings without enabling such models to approximate to a certain degree general purpose human-like intelligence and cognition. After all, an intended automatic recogniser will mostly deal with input sketches by humans, not machines. That is one vital reason why we do not directly apply image processing techniques to develop automatic recognisers that may rely on getting data from pre-compiled datasets (cf. [Arandjelović and Sezgin, 2011,Sun et al., 2012b]). Using a web-scale clipart image collection as the knowledge base of the recognition system (as in [Sun et al., 2012a] and [Sun et al., 2012b], for example) does not only dramatically increase the size needed to store the various (types of) object sketches, but also renders learning difficult, and renders online recognition time consuming. This is still correct even when recent techniques (that are developed in Big

Data analytics and Deep Learning, for example) are taken into consideration. But using an automatic recogniser that depends on generalising from a smaller set of conceptual entities (which must already be in the data base or background knowledge), we will not need all complete sketches of every object; only the specific parts that are proven to be important (i.e., the more prominently salient parts). Thus, instead of relying on millions of non-human images (on the web), fewer sketches drawn by humans, combined with a (simulation of the) generalisation process would model sketch recognition more efficiently.

Motivated by the previous discussion, we need to have an idea of how humans express their conceptual knowledge gesturally; how they outline the most important parts of their sketch-based conceptualisations of objects, and how they relate these parts when they sketch them. This should give us a practical method to approximate the realisation of sketches *à la human*.

In this article, we explain and discuss the results of two of our experiments to assess the performance of people when they deal with free-hand sketches. One experiment is concerned with the process of recognising sketched objects, while the other focusses on constructing objects via sketching. We try to empirically contribute to understanding what normal people usually prioritise when they construct sketches to represent objects, and what makes them recognise an object while it is being sketched. The focus is mainly on what makes a sketch-based feature of a specific object more salient than other features (during recognition and construction of that object's sketches). In sections 2 and 3, we try to answer the related questions, based on preliminary results of the two experiments.[1]

2 Experiment 1: Sketch Recognition

The first of our two experiments in this article aims to find "relevant parts" during the recognition process of a sketched object. In this experiment, a sketch is incrementally uncovered (while it is being drawn) by showing parts of it in a step-by-step additive way. Participants have the task to recognise as early as possible the object depicted by the (partial) sketch. Reaction times and correctness are measured in order to build hypotheses about the importance of certain parts and structural properties of a sketch for the recognition of an object. In what follows, we describe the experiment setup, procedure, results, and some analyses.

Experiment Setup: A total of 17 black-and-white free-hand sketches of miscellaneous objects (2 test trials and 15 main objects) are used for testing how each of 14 human participants recognises the underlying objects. Before conducting the experiment, the strokes that compose each sketch were pre-recorded while an artist is producing them in sequence. The recordings are played during the experiment to the participants, who are first permitted to read onscreen instructions then ask any necessary questions before beginning the experiment. The

[1] The experiments and work presented in this article are supported by the German Academic Exchange Service (Deutscher Akademischer Austausch Dienst – DAAD), as part of the $\mathtt{J_{E}SICS}$ research project (grant agreement 57247603).

participants are asked to react as soon as possible once they recognise the object being sketched, even if the played recording has not come yet to its end.

Experiment Procedure: The experiment is described using the following general steps. The participants are first trained by performing the two test trials to familiarise themselves with the experiment before they start the main trials. In a single trial of the experiment, a participant is asked to identify (or have a guessing of) an object that is being depicted. The identification should be as soon as possible (i.e., at any time during the drawing process, but as soon as the participant first recognises the object being depicted). Once the participant recognises the object, they label the object (i.e., give a name to it). In case the answer (i.e., the labelling) is correct, the participant moves on to the next object. In case the answer is incorrect, the participant is driven back to the same object at the moment they stopped at, and continue exactly with the same procedure. (N.B. this latter case would happen many times. It is repeated as long as the recording of the object being depicted is not entirely played, and the participant still gives wrong responses.) In case the depiction of the object is completed, yet cannot be recognised by the participant, the participant types "unknown" in the answer field. In both the first and second cases, the participant is given a feedback whether or not their answers are correct.

Preliminary Results and Analyses: The main aim of the experiment is to get knowledge about conceptual salience that would be enough to help a participant to correctly label the objects sketched by others. Table 1 summarises the features (or sub-concept labels) that have been found to be the most salient for 14 of the objects (main concept sketches) used in the experiment (cf. Figure 1). The left-most column of Table 1 lists the sketched objects, the middle gives a list of salient feature labels that enabled the highest percentage of the participants to recognise the corresponding object, and the third gives the percentage of the participants who were able to recognise the corresponding object using only these features. The first row in Table 1, for example, shows that 28.60% of the participants were able to identify and recognise the object "DOG" from the sub-sketch corresponding to only the features "*head*", "*ears*", and "*eyes*". Of all the participants, 71.40% were able to recognise "OWL" from the sub-sketch corresponding to only "*head*" and "*eyes*". And so on, and so forth.

Figure 1 shows the partial depictions of each of the objects that correspond to the list of concepts given in Table 1. Each of these partial depictions approximates the (visual) conceptual salience of its corresponding main concept. For example, the part of the fish *body* outlined in Figure 1 (first top row, third partial sketch from left) is enough for 57% of the participants to immediately recognise that the object being depicted would eventually be "FISH".

Note that labels of the sub-concepts, taken literally in their own or in another experiment, do not lead to a successful identification of the objects they constitute in the sketches. These literal labels need to be "seen" by the participants before their judgement, indicating that a "visual" element is playing a fundamental role in the sketch recognition process. For example, although 71.40% of

Objects (concepts):	Features (salient concepts):	Participants %:
Dog	*head / ears / eyes*	28.60
Stove	*top surface / hot plates / heat regulator*	35.70
Fish	*body / mouth / tail*	57
Owl	*head / eyes*	71.40
Palm Tree	*trunk*	50
Scissors	*finger holes / (two) blades*	28.57
Bus	*body / (one) tire*	57.14
Penguin	*head / beak / eye / body / feet*	28.57
Pineapple	*fruit / wooden skin / leaves*	50
Cactus	*stem / (two) arms*	28.57
Flower	*corolla / petals*	42.86
Giraffe	*head / ears / mouth / horns / eyes / neck*	57.14
House	*body / roof*	71.43
Umbrella	*handle / stick*	78.57

Table 1: The objects (main concepts) used in the "Sketch Recognition" experiment, their corresponding salient features (sub-concepts), and the percentage of participants who commonly recognised these features in the experiment.

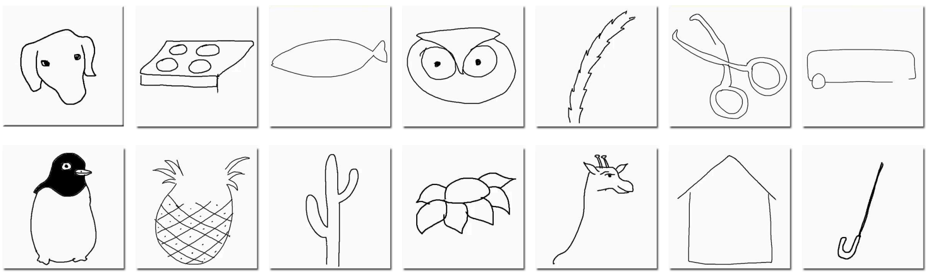

Fig. 1: Partial depictions that correspond to the most salient parts of objects of the "Sketch Recognition" experiment.

the participants were able to recognise Owl from "seeing" the *head* and *eyes* alone, the situation would be very difficult to reason that the object is Owl by letting them "know" that the object consists mainly of *head* and *eyes* as sub-concepts (or feature labels). Similarly, the three features *body*, *mouth*, and *tail*, are not enough alone to recognise that the object they compose is Fish: If you ask –not show– a person: *"what object is composed of body, mouth, and tail?"* you would probably get several answers different from Fish. Apparently, the reason for recognising Fish as the object (with *body*, *mouth*, and *tail*) while being depicted is that the strokes composing the sketch-based features are visualised, along with all their relative spatial arrangements. This observation should be kept in mind whenever considering a recognition process that depends on inputs from sensors other than merely "knowledge retrieval", because the results are not self-evident.

3 Experiment 2: Sketch Construction

When one seeks to model a cognitively inspired sketch recogniser, one at least needs an idea about the responsible cognitive processes, how they function, how they interact, and how they are utilised during sketch generation and recognition. But there is no exact brain readers that have (yet) been invented to help in translating such processes into machine understandable codes. The good news is that, there are types of human-computer interfaces that approximate (as far as cognitive science tells us) many brain activities, eye saccades, muscle or nervous system operations and signals, etc., which can help in performing the translation via simulation and modelling (cf. [Lazar et al., 2010]). For example, an eye tracker can be employed to trace the participant's gaze while sketching (which could help in inducing spatial relationships between the sketched strokes), or an electroencephalogram (EEG) may be used to record brain activities while a participant is performing an experiment, and so on.

In this experiment, we attempt one simple and direct way to inform the modeller how broadly humans conceptualise an object while they are sketching it (e.g., how some parts that constitute the object in question are always drawn, how certain parts are usually drawn in specific orders or spatial arrangement, how groups of features always come together, or how other features may or may not be drawn at all, etc.). The participants are asked to construct certain objects by outlining their shapes (i.e., sketching these objects). The resulting drawing strokes can then be recorded and analysed to extract common features among different sketches of the same object. Unlike the "Sketch Recognition" experiment, in which the several relationships between "relevant parts" of the object help to reason about its identity, the "Sketch Generation" test gives the identity and aims to construct those parts and (implicitly) indicate their relationships. This puts a particular emphasis on an object's conceptual salience, to which most people would agree while sketching this object, given the object label. This also could be used to infer spatial arrangements of- and relationships between the object's most prominently salient parts (but this particular issue is not discussed here, as it is still part of our ongoing and future work).

Experiment setup: We simply ask each participant to sketch a set of objects (main concepts), and label as many of their sketch-based features (sub-concepts) as they could. In addition to the different features they provide for each of the objects, we record the participants' drawing behaviour for the sake of further analysis (this is explained in the next paragraph). The specific subset of results in this paper[2] is based on collecting and analysing data from 8 participants working on a set with 8 preselected objects of miscellaneous categories: Fish, Bicycle, Palm Tree, Scissors, Bus, Snake, Giraffe, and House. (N.B. The two sets of participants here and of those who participated in the "Sketch Recognition" experiment in section 2 are disjoint.)

[2] Our generation experiment (on the large scale) is still ongoing, and we are still reading and analysing newer results (and for several setups).

The given results are based on sketches produced by human participants using the CogSketch environment [Forbus et al., 2008], which is only used here to allow recording the many details of each participant's sketching behaviour, such as the time taken to draw specific parts or the ordering of drawing various features. We developed an integrated Matlab code to implement both an interface and flow-controller for the entire experiment. The experimenting environment incorporates CogSkecth as an input medium within Matlab for all the processes of sketching, naming, saving, and retrieving the objects (and the sketching behaviours) for the analysis. The information collected by our code helps to achieve many proposed goals, and can later be transformed to a representation that can be understood by (and worked upon with) a variety of representation languages. In this way, the data is made ready for employing particular operations, such as generalisation, analogical, or spatial reasoning using several frameworks (e.g., an example is given in [Abdel-Fattah et al., 2015] for the case of the HDTP [Schwering et al., 2009] framework).

Experiment Procedure: The participants are asked to fill in a form, read onscreen instructions, and ask any necessary questions before beginning the experiment. There is no test trials, but there is a training phase, during which the participant gets acquainted with the experimenting environment and has time to scribble non-recorded strokes before the actual drawing –and information recording– phase starts. Each participant is then asked to sketch the preselected set of objects.

Preliminary Results and Analyses: The experiment described here seeks to get more knowledge about conceptual salience of a main concept: what are the salient sketch-based features (sub-concepts) of an object (main concept) that are commonly generated when this object is sketched? The specific aim of this experiment for the current article is twofold: to assess 'the conceptual salience commonality' and to estimate 'the conceptual salience ordering'. This means that the concern in analysing the current part of the results is both: to roughly assess the interrelationship between the sketch-based features that are *commonly generated* in drawing various objects (when the main object, which they compose, is sketched), and to roughly estimate the *ordering sequence* of an object's features by which people draw these features in average.

Regarding the *commonly generated* features of an object, it makes sense to conclude (both intuitively and empirically) that not only certain features play more essential roles than others in conceptualising the object, but they are also highly likely to be (usually) constructed when this object is sketched by anyone. This allows us to conjecture that, there will (usually) be specific strokes that are constructed while drawing the most commonly conceptualised sub-concepts of a certain main concept. These sub-concepts are the most salient sketch-based features, and seem to 'sparkle' more in the participant's mind while generating a sketch of the object depicting that main concept. Though this unproven conjecture is part of a future experiment, in which we intend to use an electroencephalogram (EEG), it is worth mentioning that it perfectly agrees with the seminal results (by Eleanor Rosch and others) that concepts exhibit *prototype effects*: the indication of membership degrees, correlating with similarity to a central member, or basic-level concept (cf. [Rosch, 1975]).

Figure 2 shows percentages of the participants who generate the various subconcepts for 4 of the main concepts of the experiment (see also Figure 3a for a fifth one). The percentages of appearance of concepts of the object BUS are given in Figure 2a, where 100% of the participants generated *wheel* and *window*, 83% generated *body* and *door*, 50% generated *light*, and so on.

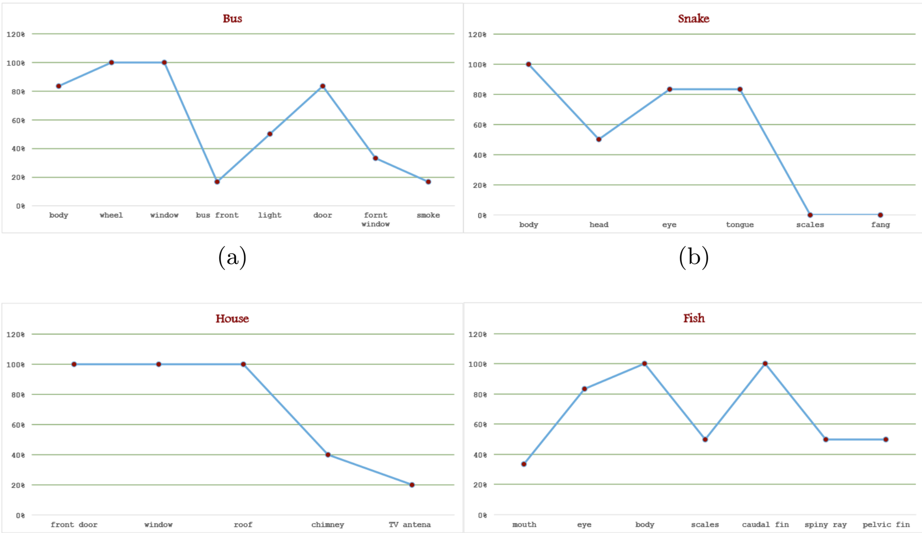

(a) (b) (c) (d)

Fig. 2: The commonly generated features for (a) BUS, (b) SNAKE, (c) HOUSE, and (d) FISH objects.

When it comes to the ordering of an object's features, the conceptual salience of features seems to be generated in sequence. By analysing the *ordering sequence*, in which the participants sketch the constituting features, the most salient features are found to come in specific order for certain objects. In our results so far, no one has ever sketched the PALM TREE's *leaves* before its *stem* (see Figure 3d). Moreover, the higher the salience of the feature the earlier it is generated in the sequence (in which the participants sketch an object). For example, the sequence in which all[3] the features of a BICYCLE object were generated is shown in Figure 3b. Among these features, *wheel* is a prominently salient feature, with 100% of the participants generating it while drawing the BICYCLE object (as Figure 3a also shows). A high majority (83%) of the participants

[3] These are "all" the features that appeared during the running of the experiment.

sketched *wheel* first, while only less than one-fifth of them (17% to be precise) sketched it fifth (and none generated *wheel* in any other order). It can also be seen from the same figure, Figure 3b, that 33% of the participants sketched *fork* as the second concept, 17% sketched it as the 3rd concept, 17% sketched it as the 4th concept, while none at all sketched it in any other order. On the other hand, *pedal* (which 50% of the participants do not generate while drawing the BICYCLE object) never comes earlier than last or next-to-last in any of the generated sequences (cf. Figure 3b). We may conclude that the most salient features are (usually) generated in order, and are (usually) generated earlier than the less salient features. (N.B. Note that it may still be the case that the simpler the feature to retrieve from memory and sketch, the earlier it is likely to be sketched.[4])

This kind of results on sketch generation is quite important for *free-hand sketch recognisers* that may need to quickly guess the object being sketched and save many computations. For instance, although *wheels* are sketched by all the participants while drawing both BUS and BICYCLE objects, *wheels* are more likely to be sketched as the first constituent for drawing BICYCLE than for drawing BUS (cf. Figures 3b and 3c). Unlike the case in sketch recognition, the sequence of the more prominently salient features in the object construction experiment seem more promising to quickly guide a successful object identification (at least to a certain degree). Maybe one would probably get BICYCLE as an answer, if one asks –not shows– a person: *"what object is composed of wheel, saddle, and spoke (in this order)?"* –maybe not; still, this needs to be justified, of course.

Concluding Notes: As the reader might have already noticed, we think that variants of the sketch construction experiment are essential for cognitively inspired models that may handle automatic sketch recognition. One needs more than the presentation in this article to elaborate on the many aspects in which this kind of experiment is useful. On the one hand, it helps in assessing each object's conceptual salience, in ranking the commonly salient features, in finding out whether a common ordering of an object's features exists (i.e., an ordering by which the participants usually draw these features), in establishing a hierarchical knowledge base of objects and their common features based on the many views of how this object can be drawn, and in hierarchically relating these objects and features together in groups of super-/sub-concept ontology (whenever interrelationships exist), among other senses of importance. On the other hand, the experiment has a variety of factors that affect not only its setup and running, but also the interpretation of results: from one's background knowledge to one's drawing skills, the result is affected; also, with several people drawing the same object, several possible, rather incomparable, views would result; and so on. Automating the recognition of, in particular, free-hand generated sketches

[4] By 'simpler sketch-based features' we mean those feature parts of the object that can be sketched using less complex drawing strokes and curves. The correlation between the 'simplicity' of sketching a feature and its 'saliency' needs to be investigated.

is a challenging task due to various causes, such as the imprecision of drawn strokes that usually constitute such sketches and the variety of ways a single object can be sketched in.

References

[Abdel-Fattah, 2014] Abdel-Fattah, A. M. H. (2014). *Utilizing Cross-Domain Cognitive Mechanisms for Modeling Aspects of Artificial General Intelligence*. PhD thesis, University of Osnabrück.

[Abdel-Fattah et al., 2015] Abdel-Fattah, A. M. H., Zakaria, W., Abdelghaffar, N., Abdelmoneim, N., Elsaadany, N., Schwering, A., and Kühnberger, K.-U. (2015). Towards the Modeling of Sketch-based Understanding: Can Computers Interpret what Non-Artists Draw? In *The Ninth International and Interdisciplinary Conference on Modeling and Using Context*.

[Andrew et al., 2009] Andrew, A., Karlson, A. K., and Brush, A. B. (2009). Investigating the use of voice and ink for mobile micronote capture. In *INTERACT 2009*. Springer Verlag.

[Arandjelović and Sezgin, 2011] Arandjelović, R. and Sezgin, T. M. (2011). Sketch recognition by fusion of temporal and image-based features. *Pattern Recognition*, 44(6):1225–1234.

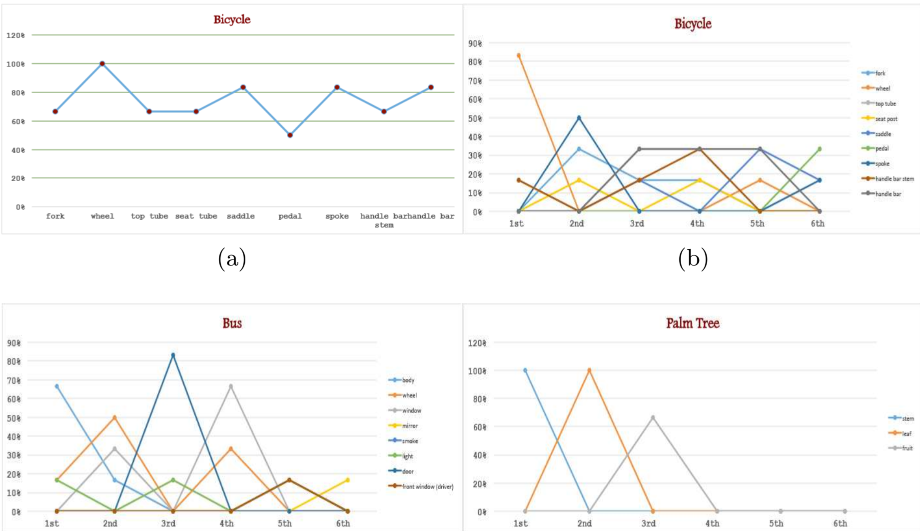

(a) (b)

(c) (d)

Fig. 3: (a) The commonly generated features for the Bicycle object. (b) The different orders in which the features of the Bicycle object are generated. (c) The different orders in which the features of the Bus object are generated. (d) The different orders in which the features of the Palm Tree object are generated.

[Csapó et al., 2015] Csapó, Á., Wersényi, G., Nagy, H., and Stockman, T. (2015). A survey of assistive technologies and applications for blind users on mobile platforms: A review and foundation for research. *Journal on Multimodal User Interfaces*, 9(4):275–286.

[Dastani and Scha, 2003] Dastani, M. and Scha, R. (2003). Languages for gestalts of line patterns. *Journal of Mathematical Psychology*, 47(4):429–449.

[Forbus et al., 2008] Forbus, K., Usher, J., Lovett, A., Lockwood, K., and Wetzel, J. (2008). Cogsketch: Open-domain sketch understanding for cognitive science research and for education. In *EUROGRAPHICS Workshop on Sketch-Based Interfaces and Modeling*.

[Gentner et al., 2001] Gentner, D., Holyoak, K. J., and Kokinov, B. N., editors (2001). *The Analogical Mind: Perspectives from Cognitive Science*. MIT Press.

[Hammond et al., 2010] Hammond, T., Logsdon, D., Peschel, J., Johnston, J., Taele, P., Wolin, A., and Paulson, B. (2010). A sketch recognition interface that recognizes hundreds of shapes in course-of-action diagrams. In *CHI '10 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '10, pages 4213–4218, New York, NY, USA. ACM.

[Lazar et al., 2010] Lazar, J., Feng, J. H., and Hochheiser, H. (2010). *Research Methods in Human-Computer Interaction*. Wiley Publishing.

[Rosch, 1975] Rosch, E. (1975). Cognitive Representations of Semantic Categories. *Journal of Experimental Psychology: General*, 104:192–223.

[Schwering et al., 2009] Schwering, A., Krumnack, U., Kühnberger, K.-U., and Gust, H. (2009). Syntactic Principles of Heuristic-Driven Theory Projection. *Cognitive Systems Research*, 10(3):251–269.

[Sciascio et al., 2002] Sciascio, E. D., Donini, F. M., and Mongiello, M. (2002). Spatial layout representation for query-by-sketch content-based image retrieval. *Pattern Recognition Letters*, 23:1599–1612.

[Sun et al., 2012a] Sun, Z., Wang, C., Zhang, L., and Zhang, L. (2012a). *Computer Vision – ECCV 2012*, volume 7572 of *Lecture Notes in Computer Science*, chapter Free Hand-Drawn Sketch Segmentation, pages 626–639. Springer.

[Sun et al., 2012b] Sun, Z., Wang, C., Zhang, L., and Zhang, L. (2012b). Query-adaptive shape topic mining for hand-drawn sketch recognition. In *ACM Multimedia*, pages 519–528. ACM.

[Wang et al., 2011] Wang, C., Zhang, J., Yang, B., and Zhang, L. (2011). Sketch2cartoon: composing cartoon images by sketching. In Candan, K. S., Panchanathan, S., Prabhakaran, B., Sundaram, H., chi Feng, W., and Sebe, N., editors, *ACM Multimedia*, pages 789–790. ACM.

[Yuan et al., 2008] Yuan, Z., Pan, H., and Zhang, L. (2008). *Advances in Blended Learning: Second Workshop on Blended Learning, WBL 2008, Jinhua, China, Augustl 20-22, 2008. Revised Selected Papers*, pages 55–64. Springer, Berlin, Heidelberg.