

IMAGING REALITY AND ABSTRACTION AN EXPLORATION OF NATURAL AND SYMBOLIC PATTERNS

Alexandra Branzan Albu¹^a, George Nagy²^b

¹Electrical and Computer Engineering, University of Victoria, Victoria, BC, Canada

²Electrical, Computer, and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY, US
aalbu@uvic.ca, nagy@ecse.rpi.edu

Keywords: Physical Scenes; Symbols; Perception and Cognition

Abstract: Understanding visual symbols is a strictly human skill, as opposed to comprehending natural scenes—which is an essential survival skill, common to many species. As an illustration of the natural vs. symbolic dichotomy, selective features are computed for differentiating a satellite photograph from a map of the same geographical region. Images of physical scenes /objects are currently captured in all parts of the electromagnetic spectrum. Symbols, whether produced by man or machine, are almost always imaged in the visible range. Although natural and symbolic images differ in many ways, there is no universal set of differentiating characteristics. With respect to the traditional branches of pattern recognition, it is tempting to suggest that statistical, neural network and genetic/evolutionary pattern recognition methods are eminently suitable for images of scenes and simple symbols, whereas structural and syntactic approaches are best for more complex, composite graphical symbols.

1 INTRODUCTION

Patterns are arrangements of perceptible elements which play a critical role in human cognition processes, such as visualization, memorization and decision-making. Furthermore, there is evidence that humans learn abstract concepts such as mathematical ones using pattern recognition techniques (Mulligan and Mitchelmore, 2009). As stated by Warren (2005), “The power of mathematics lies in relations and transformations which give rise to patterns and generalizations. Abstracting patterns is the basis of structural knowledge, the goal of mathematics learning.”


From an evolutionary viewpoint, humans first dealt with *natural patterns*, informed by their direct interactions with the environment. A small amount of “relevant” information is extracted from a large, continuous influx of data and encoded into a persistent mental structure called pattern (Del Viva, 2013). While this process is not limited to visual data (as all sensory modalities may contribute to the formation of one pattern) our paper focuses on visual patterns only. This is justified by the dominance of


the visual perception (Stokes and Biggs, 2004), as well as by the need to establish reasonable boundaries for this exploratory journey.

Natural pattern processing is a survival skill shared with other primates, allowing for generating cognitive maps of the physical environment, which encode locations of food sources, potential predators and navigation landmarks (Mattson, 2014).

Symbolic patterns are specific to humans. Symbols denote ‘something which stands for something else’ (a meaning first recorded in ‘Faerie Queene’ in 1590), thus they are representations of representations. The processing of symbolic patterns forms the basis of “unique features of the human brain including intelligence, language, imagination, invention, and the belief in imaginary entities such as ghosts and gods” (Mattson, 2014). Some simple types of symbolic patterns are embedded in our environment (for instance, traffic signs and pavement markings). Others form the basis of written language and communication (letters, digits, flowcharts, tables, etc.).

Semiotics explores the connection between signs, symbols and significance. From a semiotic

^a <https://orcid.org/0000-0001-8991-0999>

^b <https://orcid.org/0000-0002-0521-1443>

perspective, natural images fall into the category of iconic signs. Symbolic images are symbolic (arbitrary) signs, regarded as conventional and culture-specific mems for conveying concepts. These definitions do not bear directly on algorithms, features and pixels. Since natural and symbolic pattern processing do exhibit different neural mechanisms (Mattson, 2014), it seems appropriate to investigate how computer vision deals with these two meta-categories of patterns.

Among the most important shared tasks are *segmentation* and *classification*. Other common objectives are *visualization*, e.g. depth-from-shading for natural and OCR results for symbolic, and *author identification* or *counterfeit/plagiarism detection* of paintings (natural) and manuscripts (symbolic). An example of a hybrid (natural and symbolic) pattern recognition task would be a self-driving car reading all the highway signs, as well as detecting vehicles and pedestrians.

Although art critics may object, our perspective precludes attaching symbolism to a still life or an abstract painting. But Leonardo da Vinci's sketches of muscles and catapults and Edward Tufte's artful visual displays of quantitative information (in his eponymous book) *are* symbolic. Artistic applications of image processing, such as photomosaics, are skillfully explored by Tanimoto (2012).

The remainder of this paper is structured as follows. Section II presents a case study comparing a natural image with its symbolic representation. Sections III and IV discuss characteristics of natural and symbolic patterns respectively. Section V examines how pattern recognition techniques align with the natural/symbolic realms. Section VI summarizes our findings and concludes this work.

2 EXAMPLE (CASE STUDY)

Two images of the Hoover Dam area in Fig. 2.1 are chosen to illustrate the proposed dichotomy. They exemplify the potentially extreme difference between natural and symbolic images. Their respective sizes are 1050×1600 and 895×1433 pixels, thus much detail is lost in the figures.

The features below, extracted with MATLAB 2016a, show some noticeable differences, which may be quantified in many possible ways. Any ICPR participant could propose other equally plausible features. However, only experimentation on large and diverse data sets could provide statistically significant evidence that the postulated subpopulations can be objectively and accurately discriminated. Almost no

sample datasets are currently available for such experimentation (Nayef 2019). With a few exceptions (e.g. *scene text*), most of the available collections fall squarely into relatively homogenous subdomains of either Natural or Symbolic images.

Symbolic images tend to consist of high-contrast curve segments, glyphs (graphical symbols) and regions of nearly constant color because drawings and symbols have been traced for millennia using a stylus, and printing has loosely mimicked this process. Contrast helps perception; glyphs encode information. Both foreground and background of symbolic documents typically exhibit locally uniform reflectance. In natural images the distinction between foreground and background is either arbitrary, application-dependent, or refers to distance from the imaging instrument.

The logarithmic grayscale histogram (Fig. 2. 2) provides a measure of contrast. Documents usually show high peaks near opposite ends of the gray-scale, with intermediate values only at edge pixels. The proportion of edge pixels depends, of course, on the spatial sampling frequency and the point-spread function. In our map, the only two sharp peaks are near each other because the intensity of the water and land areas is almost the same. The satellite image has a wider peak for Mead Lake, and a narrow white peak due to the superimposed labels. The rest of the image has a continuous albedo distribution.

Fig. 2.3 is the 2-D Fourier transforms of the images. Strong orthogonal components are typical of document images because of their rectilinear print layout, but much less so of line drawings and maps. The small higher-frequency components of our map fall outside the range of FFT coefficients visible in the figure.

In addition to extracting the intensity distribution, we have chosen for this illustration features that are sensitive to local variability like edges. Fig. 2.4 indicates that the natural image has a far greater density of Canny edge features than the symbolic image (Canny, 1986). Although their sizes differ by only 30%, the edge count is 251,332 vs. 14, 320. The superimposed geodetics are detected in the satellite image, and most of road network, barely visible in Fig. 2.1, on the map. The FAST features of Fig. 2.5, extracted with the popular algorithm proposed by Rosten and Drummond (2005), exhibit a similar configuration (8316 vs. 353 corners).

None of the above features suffice by themselves for differentiating natural from symbolic images. For example, the snow-covered shores of unfrozen lakes offer high contrast like printed pages, fingerprints abound in curvilinear features as do caricatures, and

both documents (*symbolic*) and aerial photos of cities (*natural*), exhibit a profusion of corners and edges. Furthermore, many of the above features are class-conditionally statistically dependent. Automated classification would require many more features and a highly nonlinear classifier.




Figure 2.1: Examples of a natural and of a symbolic image. Source: <https://www.lakemeadcruises.com/discover/area-maps/getting-here/>




Figure 2.2: Logarithmic intensity histogram.




Figure 2.3: 2-D Fast Fourier Transforms (FFTs)




Figure 2.4 : Canny edges




Figure 2.5 FAST features


3 NATURAL VISUAL PATTERNS

Fig. 3.1 is a montage of photographs that often appear in image processing research and illustrate the variety of aspects (color, contrast, level of detail) that affect processing. The source site, ImageProcessing-Place.com, also offers free downloads of many larger image collections. Although Fig. 3.1 is drawn from the visible regions of the spectrum, natural images span almost twenty orders of magnitude in wavelength or frequency (Fig. 3.2). Regardless of their source spectrum, they can be rendered to be visible at an appropriate scale for human inspection.



Figure 3.1: Standard (natural) test images http://www.imageprocessingplace.com/root_files_V3/image_databases.htm

Photography gained mass appeal soon after its invention early in the 19th Century. It became ubiquitous after digital cameras were grafted onto cell phone and photo-sharing social media applications, such as Facebook and Instagram, gained wide popularity. An early quantitative application was cartography with photographs from balloons (Redmond, retrieved 2020). Natural Image Processing (IP) started in the 1950's with the analysis of photographs of the tracks of elementary particles in spark, bubble and cloud chambers. Algorithmic path tracking was a disruptive technology, as it displaced the dozens of operators who had traced the tracks on projection screens. Computer Cartography and Geographic Information Systems (now Geospatial Data Processing) eventually grew to encompass earth observation and weather satellites that currently produce over one million images per day. Many earth and ocean observation facilities produce a huge amount of visual data, which exhibits typical Big Data problems (storage, curation, provenance, manipulation).



Cosmic-ray muography (muon tomography)
 Atomic force and electron microscopy
 Medical and industrial radiography (X-rays)
 Industrial surface inspection, fluorography (UV)
 Photography, microscopy, telescope (visible light)
 Night vision, thermography, FLIR, LIDAR (IR)
 Weather, traffic and military RADAR (microwaves)
 Radio-telescope, MRI (radio frequency)
 Medical and industrial ultrasound

Figure 3.2: Natural images span the entire electromagnetic and sound spectra. <https://www.britannica.com/science/electromagnetic-spectrum>

Modalities outside the narrow visible range reveal details varying in scale from nanometers for atomic lattices, to micrometers for biological cells, at “human scale” for animals, plants and organs, and light years for astronomical observations. A single infrared, visible or radar image for automotive, ship and aircraft applications may cover an area with a diameter of a few meters or hundreds of kilometers. Thus some Natural IP modalities have extended our understanding of the world far beyond the original goals inherited from animal vision of wayfinding and navigation in environments constrained by our limited visual abilities. Moreover, some Natural IP techniques allow us to gain not only structural, but also phenomenological insights. Many striking examples of such techniques come from medical imaging, where modalities such as computed tomography (CT), magnetic resonance imaging (MRI), Doppler ultrasound, scintigraphy, single photon emission computed tomography (SPECT) and positron emission tomography (PET), rendered in the visible spectrum, are used to examine physiological and metabolic phenomena.

Scientific and industrial Natural IP often includes input from non-imaging sensors. Furthermore, the most interesting applications require processing groups of images. The grouping may be based on spatial contiguity (mosaicking or slices of a 3-D volume), sparse time sequences (monitoring the growth of vegetation or beach erosion), or movie-rate sequences (motion from video). Natural IP now includes 2 1/2 D, 3D, and 4D, grayscale, color, and multispectral images. For example, sequences of high energy X-rays (from a synchrotron) have been used to study crack propagation in concrete under load (Landis et al., 2007). Cosmic-ray muography, first used to map hidden chambers in pyramids (Alvarez et al., 1970), is used for inspecting nuclear waste sites (Linkeos Technology Ltd, 2020).

Current applications include video from multiple cameras for analyzing traffic, athletic events, and crowd activity in premises with high security concerns. Natural IP is gradually merging into Computer Vision because images from robots, drones, self-driving cars and wearable cameras must accommodate variable lighting and relative motion between multiple sensors and targets. As we will see in the next section, Document Analysis is moving in an entirely different direction, shifting emphasis from images to computer-native text and graphics.

4 SYMBOLIC VISUAL PATTERNS

Most symbolic images are, by definition, *documents* (or parts of documents). This overarching category includes books, magazines, newspapers, handwritten letters and notes, plans and diagrams, musical scores, tables, maps, charts and graphs.

The first patents on Optical Character Recognition (OCR) were filed more than one hundred years ago, but until the 1960s OCR had to run on hardwired machines because computers took a long time to process even a 256 x 256 image. Figure 4.1 shows postage stamps from the CCITT test sequence prepared for the standardization of facsimile in the 1970's. They are available full-size at the website in the figure caption (of the International Telecommunications Union) which also houses many excellent sets of test data and calibration charts with complete metadata. Some (like the graph and the circuit diagram) may have been originally intended for *visualization*.

Many applications that fueled optical character recognition and document image processing in the last century have virtually disappeared (Nagy, 2016). The list includes postal address reading, bank check

reading, and invoice image processing. Forensic document analysis is giving way to white-hat hacking (Al-Muhammed and Daraiseh, 2018). Research on document analysis is shifting from processing images to manipulating documents already in a computer-readable symbolic format such as plain text or XM. Current objectives include deep document understanding, search, summarization, translation, information extraction, table analysis, and sketch understanding (Nagy, 2016).




Figure 4.2: Symbolic test images for evaluating compression algorithms for facsimile transmission. <https://www.itu.int/net/itut/sigdb/genimage/test24.htm>

Until the invention of the printing press, every image could be traced to the person or group who prepared (or copied) it. The largest libraries contained only a few thousand items. With the advent of printing and lithography, the direct connection between the image and its creator was lost. The number of physical images grew exponentially because each reproduction could be replicated at will. By the end of the last century, public and university libraries (and museums) had to store much of their holdings at remote locations. Progress in digitization and storage technology is now impelling libraries to move their shelves to a Cloud.

We are now on the threshold of losing even the indirect link from image to creator (author, printer,

artist, draftsman, or composer). Computers can dissect and reassemble symbolic images in myriad ways. The provenance of the whole or of parts thereof becomes untraceable. Furthermore, computers routinely convert signals from measuring instruments into symbolic images. More and more symbolic images have no human genesis, which is worrisome since computer-generated semantics may be arbitrary and not consistent with human reasoning, values, and responsibilities.

5. NATURAL VS SYMBOLIC PATTERN RECOGNITION

While recognizing and categorizing patterns is an essential philosophical endeavor first formulated by Aristotle (Ammonius, 1991), from a more pragmatic viewpoint it can be considered as a critical step in decision-making. Natural and symbolic visual patterns support, in general, different decision paths. For instance, recognizing landscape cues supports wayfinding, while recognizing written words supports reading and comprehension. A question arises naturally: for choosing the ‘right’ pattern recognition technique, does it matter whether patterns depict some aspect of the natural world or if they belong to a more abstract (symbolic) realm?

The design of a computational technique for recognizing visual patterns may start by contemplating two interrelated questions:

- a) what type of data representation is most relevant for describing the patterns of interest?
- b) what formalisms and methodologies are associated to the data representation?

The answers to these questions lay the foundations of three main schools of thought (statistical, structural, and syntactic) in computer-based pattern recognition (which might or might not be inspired by biological mechanisms). The underlying principles are compared below. The interested reader is referred to Bunke and Riesen (2012) for more details regarding this comparison.

Statistical pattern recognition represents a given pattern by a *feature vector* of fixed length n (i.e., as a point in an n -dimensional feature space) which enables the use of a rich arsenal of algorithmic tools grounded in linear algebra and probability. However, representing patterns via a simple concatenation of features has two main limitations, namely:

- a) the fixed length constraint, which prevents tailoring the representation to the complexity of the pattern;
- b) the difficulty of encoding binary or higher-order relationships that may exist between different components of the pattern.

Both limitations are elegantly addressed by *structural and syntactic pattern recognition*, which encode patterns using sophisticated paradigms. *Structural techniques* are intrinsically associated with graph-based representations, which allow for describing patterns by decomposing them into semantically meaningful parts (primitives) and describing properties of these parts as node labels, and inter-part relationships as edges. An illustrative example of a graph-based approach for architectural symbol recognition is provided by Lladós and Sánchez (2003).

The *syntactic approach* to pattern recognition is inspired by formal language theory, and attempts to describe a complex pattern by decomposing it into a set of smaller, simpler patterns, which are connected via grammatical rules (Searls and Taylor, 1992). It is thus similar to the *structural approach*, but it is less popular because of the difficulty of defining grammars for parsing visual entities. Some successes have been reported in early works such as (O’Gorman, 1988) and (Ripley and Hjort, 1995).

Structural and syntactic pattern recognition techniques rely upon rich, complex representations. This is both a blessing and a curse, since there is little mathematical structure to support the analysis of such representations. It becomes thus obvious that statistical and structural pattern recognition techniques exhibit complementary strengths and weaknesses, which has motivated research on combining data-rich, structural representations with statistical analysis tools (Bunke and Riesen, 2012).

The deep learning revolution, occurring within the last decade, has clearly established the dominance of the statistical school of thought over the other two. Can deep learning be considered as belonging to statistical pattern recognition? A positive, mathematically justified answer to this question is offered by Ripley and Hjort (1995), who outline two main paradigms (sampling and diagnostic) for learning posterior probabilities in pattern recognition. Core to the deep learning paradigm is the concept of *neural networks*, which can be thought as a generalization of the diagnostic paradigm. This paradigm learns posterior probabilities directly from examples in the training set which are similar to the sample to be recognized. Bishop (2006) also

considers neural networks as efficient models for statistical pattern recognition, as they provide a convenient solution to the curse of dimensionality (Bellman, 1961). This solution formulates the non-linear mapping function (from the feature space to the classification space) as a linear combination of non-linear activation functions (the ‘neurons’).

The appeal of deep learning techniques might be partially explained by the elimination of the feature extraction step from traditional statistical pattern recognition pipelines (i.e., the handcrafting process) which involved a careful analysis of the dimensions of variability of the patterns of interest, as well a study of visual cues used by humans for performing the same detection/localization/classification task. Deep learning networks accept image patches as inputs, and discover not only the mapping from the feature vector representation to the output, but also the representation itself; thus, they perform *representation learning* (Goodfellow et al., 2016). This works well for most natural patterns where image patches of reasonable size are information-rich. However, some symbolic patterns may consist of just a few linear/circular segments; this is the case of symbols composing architectural floor plans (Rezvanifar et al., 2019). In such cases, examples from small-sized training datasets do not provide enough information for a successful learning process. Yosinski et al. (2014) show that transfer learning procedures, which allow to learn general features from base networks trained on rich, generic datasets, and specific features from target networks and smaller datasets, yield decreased performance when the distance between the base task and the target task increases. This explains, in part, the limited success of deep learning methods on sparse symbolic patterns.

We cannot ignore game-changing results of deep learning architectures on both natural and symbolic public datasets (Farabet et al., 2013; LeCun et al., 1998; LeCun et al., 2015). Deep learning networks presumably also played a role in digitizing Google Books, the largest collection of symbols in the word. However, none of the applications supported by these public datasets suffer from sparsity of data, such as the one in (Rezvanifar et al., 2019). The plethora of labeled and unlabeled training data (sometimes millions of samples) overcomes any benefit of syntactic or structural representation of human insights.

To summarize, statistical methods work better for most natural patterns and simple symbolic patterns (such as digits/printed or handwritten characters or musical scores), while structural/syntactic techniques

are more suitable for more complex and/or sparse symbolic patterns. This is, of course, only a starting point, which is nevertheless useful for pondering the entire pattern recognition repertoire before delving into a more nuanced exploration.

Indeed, the boundaries of structural and statistical approaches are blurring (e.g. probabilistic grammars, Markov random fields). The recent emergence of Graph Neural Networks which inject deep learning into computational graph analysis is of particular interest (Renton et al., 2009; Battaglia et al., 2018). Nevertheless, applications (such as multimedia and document image analysis) relying heavily on symbolic patterns are still mentioned as belonging to the area of structural approaches, as shown in the 2020 Call for Papers of the S+SSPR Workshop <https://www.dais.unive.it/sspr2020/call-for-papers/>.

6. CONCLUSIONS

We explored some characteristics that can reveal whether the source of an image is a real-world scene or an abstract concept. The proposed distinction between natural and symbolic images focuses attention on an essential difference between human and animal cognition and suggests a pathway to advance the study of both. The distinction also helps explain why syntactic and structural methods are seldom applied to scenery and to scientific imaging (especially beyond the visible spectrum), and the popularity of statistical and neural network approaches wherever human annotation becomes overwhelming. The scarcity of databases of heterogeneous (symbolic AND natural) image samples confirms our intuition regarding the fundamental nature of the distinction.

Our future work will address the differences *within* both natural and symbolic images; we intend to survey image types in computer vision and image processing literature, which will hopefully clarify their links to pattern recognition methods.

We also intend to explore interactions and mappings between natural and symbolic patterns. The world of visual art (left out of this preliminary exploration) offers abundant opportunities for studying how natural scenes are mapped onto more abstract, symbolic representations. Augmented reality environments will enable us to look at symbiotic co-occurrences of natural and symbolic patterns.

REFERENCES

- Al-Muhammed, M.M.J., Daraiseh, A., 2018. On the Fly Access Request Authentication: Two-Layer Password-Based Access Control Systems for Securing Information. *Int. J. Information Technology* vol. 8: 6, pp. 2598-2611.
- Alvarez, L. W., et al., 1970, Search for hidden chambers in the pyramids,” *Science*, vol. 167:3919, pp. 832-839.
- Ammonius, 1991. *On Aristotle Categories*, S.M. Cohen and G.B. Matthews, (trans.), London/Ithaca: New York.
- Battaglia, P. et al., 2018. Relational inductive biases, deep learning, and graph networks, *arXiv:1806.01261v3*.
- Bellman, R., 1961. *Adaptive Control Processes: A Guided Tour*, Princeton University Press.
- Bishop, C.M., 2006. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg.
- Bunke, H., Riesen, K., 2012. Towards the Unification of Structural and Statistical Pattern Recognition. *Pattern Recognition Letters*, 33: 811-825.
- Linkeos Technology Ltd, 2020. Cosmic-ray muography, <https://www.lynkeos.co.uk/>.
- Canny, J., 1986. A Computational Approach to Edge Detection. *IEEE Trans. on Pattern Anal. Machine Intell.*, vol. 8:6, pp. 679-698.
- Del Viva, M., Punzi, G., Benedetti, D., 2013. Information and Perception of Meaningful Patterns. *PLoS ONE* vol. 8:7.
- Farabet, C., Couprie, C., Najman, L. and LeCun, Y., 2013. Learning Hierarchical Features for Scene Labeling. *IEEE Trans. on Pattern Anal. and Machine Intell.*, vol. 8: 35, pp. 1915-1929.
- Fu, K.S., 1982. *Syntactic Pattern Recognition and Applications*. Prentice-Hall, 1982.
- Goodfellow, I., Bengio, Y., Courville, A., 2016. *Deep Learning*, MIT Press.
- Landis, E.N., Zhang, T., Nagy, E.N., Nagy, G., Franklin, W.R., 2007. Cracking, damage and fracture in four dimensions. *Materials and Structures*, vol. 40:4, pp. 57-364.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep Learning. *Nature* 521, 7553, pp. 436-444.
- LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., 1998. Gradient-based learning applied to document recognition, *Proceedings of the IEEE*, vol. 86:11, pp. 2278-2324.
- Lladós, J., Sánchez, G., 2003. Symbol recognition using graphs. *Proc. of IEEE Int. Conf. on Image Proc. (ICIP)*, pp. 49–52.
- Mattson, M., 2014. Superior Pattern Processing is the Essence of the Evolved Human Brain. *Frontiers in Neuroscience*, vol. 8: 265.

Mulligan, M., Mitchelmore, M., 2009. Awareness of pattern and structure in early mathematical development. *Math. Educ. Research Journal*, vol. 21:2, 33-49.

Nagy, G., 2016. Disruptive developments in document recognition. *Pattern Recognition Letters* vol. 79, pp. 106-112.

Nayef, N, Patel, Y, Busta, M, Chowdhury, P.N., Karatzas, D., Khlif, W., Matsas, J., Pal, U., Burie, J-C, Liu, C.N., Ogier, J.M., 2019. ICDAR2019 Robust Reading Challenge on Multi-lingual Scene Text Detection and Recognition-- RRC-MLT-2019, arXiv:1907.00945 [cs.CV]

O’Gorman, L., 1988. Primitives Chain Code. *Proc. of IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 1988.

Redmond, E., 2020. Places in Civil War History: Aerial Reconnaissance and Map Marketing, retrieved from <https://blogs.loc.gov/maps/2017/08/places-in-civil-war-history-aerial-reconnaissance-and-map-marketing/>

Renton, G., Heroux, P., Gauzere, B., Adam, S., 2009. Graph neural network for symbol detection on document images. *Proc. Int. Conf. Doc. Anal. Recognit. Workshops (ICDARW’09)* vol. 1, pp. 62–67.

Rezvanifar, A., Cote, M, Branzan-Albu, A., 2019. Symbol spotting for architectural drawings: state-of-the-art and new industry-driven developments. *IPSJ T Comput Vis Appl* vol. 11: 2.

Ripley, B.D., Hjort, N.L., 1995. *Pattern Recognition and Neural Networks* (1st. ed.). Cambridge University Press.

Rosten E., Drummond, T., 2005. Fusing Points and Lines for High Performance Tracking. *Proc. of 10th IEEE Int. Conf. on Comp. Vision (ICCV)*, pp. 1508-1515.

Searls, D.B., Taylor, S.L., 1982. Document Image Analysis Using Logic-Grammar-Based Syntactic Pattern Recognition. In: Baird H.S., Bunke H., Yamamoto K. (eds) *Structured Document Image Analysis*, Springer.

Stokes, D., Biggs, S., 2014. *Dominance of the Visual*. Oxford Scholarship.

Warren, E., 2005. Young children’s ability to generalize the pattern rule for growing patterns. *Proc. of the 29th Conf. of the Int. Group for the Psychology of Math. Educ.*, vol. 4, pp. 305-312.