



Coherence in typeface design: visual similarity of characters in Cyrillic, Devanagari, and Latin

David Březina
September 2018

Thesis submitted for the degree of Doctor of Philosophy
Department of Typography & Graphic Communication

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

David Březina

This research project was funded by the Arts & Humanities Research Council under the auspices of the Design Star Doctoral Training Centre.

Thesis typed and typeset by the author in typefaces Yrsa, Skolar Sans, and Input.

To mum, dad, and Kuba.
And to Anna who supported me
all the way, although I made it hard.

Acknowledgements

I would like to thank Dr Mary C. Dyson, Dr Fiona G. E. Ross, and Keith Chi-Hang Tam for being the best supervisors I could ask for.

Also, I will always remember the hospitality and encouragement of the ANRT, Thomas Huot-Marchand and Bart Lamiroy in particular.

To Andrea Churchill Wong, Chris Earl, Maxim Zhukov, Maria Doreuli, Nils Frahm, Kiasmos, and my PhD colleagues: thank you for all your help and moral support.

Abstract

This thesis explores the visual similarity that underlies the coherence in the design of individual typefaces. Typeface designers aim to achieve a unifying coherence in their typefaces, so that characters can be identified individually as well as belonging together giving rise to an overall style. The objective is to determine whether the coherence perceived by readers differs from the coherence intended by designers. The research is cross-disciplinary, combining empirical studies of readers' perceptions with a computational model that is based on relevant typeface design knowledge.

Character similarity is studied in multiple different typefaces (fonts) intended for continuous reading in Cyrillic, Devanagari, and Latin scripts. The studies were conducted online to collect a large number of responses. The participants were presented with a sequence of character triplets. They were asked to identify the odd one out in each of these triplets judging by their visual similarity, thus making a statement about the similarity of the two complementary characters. This method studies the similarity in context, which provides more refined details about participants' similarity judgements.

The model interprets characters using two kinds of features: more specific parts and more general roles. The model learns the relative saliences of these features from a subset of the data collected in the studies. This allows the model to predict participants' responses to the triplets from the studies and for other, unseen triplets. Additionally, the model can provide explanations of the criteria participants used in their similarity judgements and can generate similarity matrices.

The model achieved high scores when predicting response probabilities and identifying the overall odd ones out. A view of coherence that is supported by readers' perception can be used to assist designers in their creative process, help with fonts' quality assessments, and contribute to readability research and multi-script typography.

Contents

1. Introduction 8

- 1.1. Research questions 11
- 1.2. Method 12
- 1.3. Impact 14
- 1.4. A few notes about the scope and style of this thesis 17
- 1.5. Note on notation and figures 17

2. Theory and definitions 18

- 2.1. Character 18
- 2.2. Character concept 20
- 2.3. Writing and script 25
- 2.4. Typeface 26
- 2.5. Similarity and coherence 28
- 2.6. Style 29

3. Overview of the scripts studied 31

- 3.1. Latin 32
- 3.2. Cyrillic 36
- 3.3. Devanagari 40
- 3.4. Representative characters 46

4. Approaches to character conceptualization 48

- 4.1. Shapes as contours 49
- 4.2. Shapes as strokes 51
- 4.3. Transformation of shapes 56
- 4.4. Composition of shapes 58
- 4.5. Generalization of shapes 66
 - 4.5.1. Typographic attributes 66
 - 4.5.2. The role hypothesis 67
- 4.6. Towards modelling of coherence 70

5. Studies with readers 71

- 5.1. Asking about similarity 71
 - 5.1.1. Enquiry about parts 72
 - 5.1.2. Enquiry about similarity 74
- 5.2. The approach taken 76
 - 5.2.1. Triplet sequence and character pair similarity 77
 - 5.2.2. Presentation 78
- 5.3. Organisation of the studies 81
 - 5.3.1. Questionnaire 81
 - 5.3.2. Triplet trials 84
 - 5.3.3. Language variants 86
- 5.4. Scope 86
 - 5.4.1. Typefaces 86
 - 5.4.2. Characters and triplets 87
 - 5.4.3. Kinds of triplets included 95
- 5.5. Technical implementation 96
- 5.6. Pilot study 97
- 5.7. Participants 97
 - 5.7.1. Recruitment 97
 - 5.7.2. Conversion rate and dropout 97
 - 5.7.3. Devices used 98
 - 5.7.4. Comments from participants 98
- 5.8. Pre-processing 99
- 5.9. Data set overview 100
- 5.10. Participants' background data 102
- 5.11. Response data 106
 - 5.11.1. Density of the data 106
 - 5.11.2. Response frequencies 107
 - 5.11.3. Participants' agreement and triplet difficulty 108
- 5.12. Specific character pair similarity measure 111
- 5.13. Generalized character pair similarity 113
- 5.14. Character distinctiveness and perceivability 119
- 5.15. Comparison of responses 120
 - 5.15.1. Comparison of repeated triplets between typefaces 120
 - 5.15.2. Agreement between groups of participants 123
 - 5.15.3. Comparison between experts and non-experts 124
 - 5.15.4. Comparison between native and non-native, and fluent and non-fluent participants 128
- 5.16. Visual vs. linguistic interpretation 132
- 5.17. Possible methodological improvements 134
- 5.18. The use of the data 135

6. Modelling character similarity and coherence 136

- 6.1. Modelling approaches 137
- 6.2. The proposed model of coherence 145
- 6.3. Character conceptualization 149
 - 6.3.1. Saliency 152
 - 6.3.2. Parts 152
 - 6.3.3. Roles 161
- 6.4. Concepts' similarity 167
 - 6.4.1. Generalized approach 167
 - 6.4.2. Contextual approach 169
- 6.5. Model training 171
- 6.6. Technical implementation 172
 - 6.6.1. File storage 172
 - 6.6.2. Data structures 174
 - 6.6.3. Calculating predictions 175
- 6.7. Model assessment 179
 - 6.7.1. Convergence of the model 181
 - 6.7.2. Cross-validation 183
- 6.8. Predictions and explanations 187

7. Conclusion 193

- 7.1. The theory 195
- 7.2. Study format and method 196
- 7.3. Data set 197
- 7.4. Discussion 197
- 7.5. Harmonious and essential future prospects 198

Bibliography 199

Appendix 1: character shapes studied 213

Appendix 2: observations and predictions 228

Appendix 3: similarity matrices 334

Appendix 4: cluster plots 340

Appendix 5: distinctiveness measures 346

Appendix 6: typeface comparisons 350

Appendix 7: concept diagrams 372

Appendix 8: typeface mappings 405

Appendix 9: Python code 415

Appendix 10: similarity matrices compared 443

1. Introduction

Whenever there is a series of products, be it furniture, icons, books, or letters, designers aim to achieve some level of unifying coherence among the individual items in the series. In typeface design, one of the important qualities of professional typefaces is that the shapes of characters, although they can be identified individually, can also be perceived as belonging to a single set. This fundamental attribute of quality typefaces has already been mentioned in early treatises dedicated to the craft, such as those by Joseph Moxon (1703) and Pierre-Simon Fournier le jeune (Fournier, 1706; Carter, 1995). It was mentioned more recently by Fernand Baudin, who stated: “For the graphic designer and the typographer, [the alphabet] is a system of signs with a visual logic all of its own.” (Baudin, 1989, p. 23). The ability to create a sense of coherence across different character shapes is an essential skill that typeface designers need to acquire together with the ability to create characters that are legible and readable.¹

Besides adhering to tradition rooted in handwriting and other forms of manual lettering, there are at least three motivations for designers to create typefaces with a visual coherence. The first one is the association between coherence and high-quality typefaces already mentioned above, i.e. that coherent typefaces are seen as being of quality. Note that whether coherence constitutes a positive aesthetic value depends on the worldview of a particular typeface designer and contemporary design trends. The second motivation seems to be user-centric. In any typographic design, typefaces are important building blocks contributing towards the overall style of a document. In the complex environment of typographic design, they become tools, and need to function predictably so that the designers who use them have firm control over the aesthetics and atmosphere of the whole work. For example, it would be a challenge for a graphic designer to use a typeface that would sometimes appear formal and other times informal, depending on the text it is used to set. The third motivation is related to effectiveness of production. The process of designing a typeface can be very tedious since it typically involves drawing and redrawing hundreds of character shapes. Occasional reuse of existing shapes and sub-shapes within a typeface helps to reduce the workload. Such reuse may involve sub-shapes being copied verbatim from one character to another, but in professional text typefaces this typically involves subtle modification to adapt the copied sub-shapes to a new visual context. In this respect coherence can be seen as a side-effect of striving for increased effectiveness.

1. The definitions of legibility and readability used here are by Walter Tracy (2003, p. 31). *Legibility* refers to the ease of recognition of character shapes while *readability* refers to the ease of comprehension of texts or documents set using character shapes produced using a particular typeface.

This does not necessarily mean that all typefaces have to be as coherent as possible. Sometimes, incoherence is called for by the designer or the project brief. It is simply to say that there are certain perceived benefits to coherent typefaces. Furthermore, it should be noted that it remains to be reliably determined whether coherent typefaces are easier to read. Psychological studies dealing with so-called font tuning (Sanocki, 1987; Sanocki, 1988; Sanocki, 1992; Gauthier et al., 2006; Walker, 2008; Sanocki & Dyson, 2012) provide evidence suggesting that there might be a correlation between typefaces' uniformity and speed of reading. Experiments attempting to link coherence and legibility or readability, however, require a clear, formal description of coherence in the first place. This research does not try to question the motivations described above. Instead, it attempts to analyse how typeface coherence is realized, and provide a formal description so that more productive discussions about the utility of coherence can be had.

In this work, coherence is understood as an overall quality that refers to a group of shapes. Similarity, on the other hand, refers to relationships between individual shapes in pairs in particular. The main challenge is to provide a detailed description of similarity judgements involved in perceived coherence. Rather than stating that letters a, b, d look coherent, the research aims to clarify what it is that makes the individual character pairs a, b, b, d, and a, d look similar in this context, to what extent they are similar, and what this means for the overall level of coherence. It is important to emphasize that there can be different degrees of coherence (see figure 1.1). Moreover, the term coherence in this work refers specifically to similarity relationships among character shapes from a single typeface. Coherence and similarity among characters from different typefaces is not studied.



FIGURE 1.1: The degree of apparent coherence among the character shapes in each line decreases from top to bottom. The effect is caused by deliberate mixing of different typefaces in each line. The first line uses shapes from a single typeface while the second and third lines mix multiple different typefaces. The differences in the second line are subtle and hard to spot at first. The third line also includes shapes with different stroke thicknesses which makes spotting incoherences easier.

To be applicable in practice, such descriptions should be relatable to typeface design processes. The aim is to also describe coherence in a general way that would work for more than one typeface or script. Therefore, such a description should be as independent of a particular technology as possible. This also suggests that the viewpoint from which coherence is studied should be that of the human, not of the machine. The research considers the reader's perspective and the way readers perceive and interpret the appearance of characters. The typefaces are studied as perceived by experienced readers and not, for example, by beginner readers.

The aspiration for a general theory is also a reason for conducting the work with multiple scripts. Three scripts – Cyrillic, Devanagari, and Latin (see figure 1.2) – have been selected for analysis in order to set up a challenging domain for the theory: Latin as a script for which numerous character representations and typographic systems have been devised; Cyrillic as a script that bears partial resemblance to Latin, yet remains significantly different in some respects; and Devanagari as a script clearly distinct from Latin and Cyrillic, both visually and linguistically, as Devanagari is a syllabic script.

Кириллица: абвгдежзи
 देवनागरी: कखगघङचछजझ
 Latin: abcdefghijklmno

FIGURE 1.2: From top to bottom: Cyrillic, Devanagari, and Latin scripts in typefaces Georgia (Cyrillic and Latin) and Murty Hindi (Devanagari).

ABCDEFGHIJabcdefghij
 ABCDEFGHabcdefghij
 ABCDEFGHabcdefghij

FIGURE 1.3: Common text typefaces represent more conventional designs: Times New Roman (top), PT Serif (middle), and Verdana (bottom).

The focus is on text typefaces, that is, typefaces that are generally recommended for longer, continuous reading, usually set in sizes around 9 to 14 points and used, for example, for longer articles, books, newspapers, or theses, printed as well as on the screen. The main reason for this choice is that text typefaces are considered to be conventional interpretations of the script they represent. Therefore, there is a greater chance that their study will contribute more towards fundamental ideas about scripts than would a study of unconventional designs (compare figures 1.3 and 1.4).

Notably, the research is not attempting to create a new way of designing typefaces or a model of character recognition. The aim is to describe the coherence in existing typefaces and draw more general conclusions about conventional relationships between characters within the selected scripts. These relationships will be referred to as *script conventions* in this work.



FIGURE 1.4: More unconventional typefaces: Alcoholica (top), Sutturah (middle), and Anglaise (bottom).

1.1. Research questions

The aims of this research can be summed up by the following research questions:

- i. What are the principles underlying the coherence in design of individual typefaces as perceived by readers?
- ii. Can a single theory describe perceived coherence in typeface design in different scripts? In other words, is it possible to develop means of commenting on character similarity relationships in typeface design for more than one script?
- iii. Does the coherence perceived by readers differ from the coherence intended by designers?

1.2. Method

To explain typeface coherence, i.e. coherence of character shapes produced using the same typeface, one has first to explain how character shape similarity works. The study of similarity between character shapes is therefore central to this enquiry. The research consists of two main components:

- i. testing with readers as participants who provide data about their perception of similarity,
- ii. a theoretical, computational approach to similarity and coherence related to expert knowledge about character shape relationships and construction.

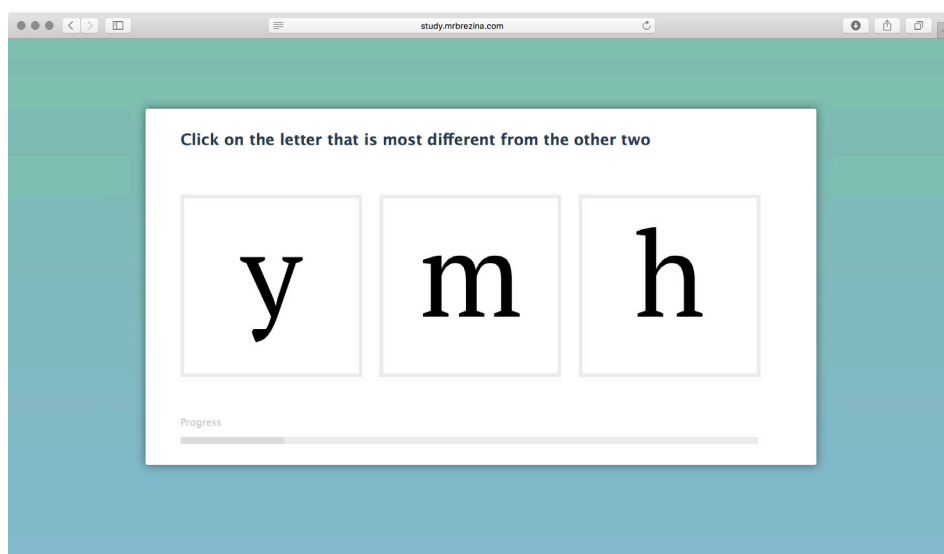


FIGURE 1.5: Participants in the studies were asked to select the most different character (the odd one out) from each in a series of character triplets.

The first component, testing with participants, is further discussed in chapter 5: *Studies with readers*. By way of introduction it should suffice to say that in studies conducted online, experienced readers performed a series of simple discrimination tasks on triplets of character shapes produced using the same typeface (see figure 1.5). For each triplet, e.g. y, m, h, they were asked to identify the character they found visually most different from the other two (the odd one out). The collected data has been analysed and used to derive response frequencies describing perceived similarity of character pairs within the context of the individual triplets. This approach to testing allowed the presentation of the character shapes without any degradation and at a comfortable size.

The second component involves a construction of a theoretical, computational model of coherence based on expert knowledge. Characters are interpreted using mental concepts and techniques typeface designers use when creating typefaces. Character shapes are formally conceptualized in terms

of their features: parts, which are more specific, and roles, which are more general. The resulting concepts are included in the computational model and used to calculate predictions regarding participants' responses from the study. Assessment of the model is by extension also a test that investigates to what extent expert knowledge can predict the observed similarity judgements of experienced readers. The typeface-design and typographic expert knowledge is further discussed in chapter 4: *Approaches to character conceptualization* while the approach to modelling and the assessment of the resulting model is detailed in chapter 6: *Modelling character similarity and coherence*. The overall research design is in figure 1.6.

It is important to consider that the model is developed independently of the designers of the studied typefaces. As such, it is only an approximation of the designers' original intentions. Umberto Eco (2015) distinguishes between the intention of the author, i.e. the typeface designer, the intention of the reader, and the intention of the text, i.e. the typeface in this context. In this sense, the results of studies with participants represent the perspective of the reader while the theoretical model stands between the perspective of the designers and that of the typefaces themselves.

As will be discussed in the conclusion, aggregation of the data collected in the studies and the generalisation of this data using the model can inform our understanding of coherence in the typefaces studied as well as the scripts' conventions.

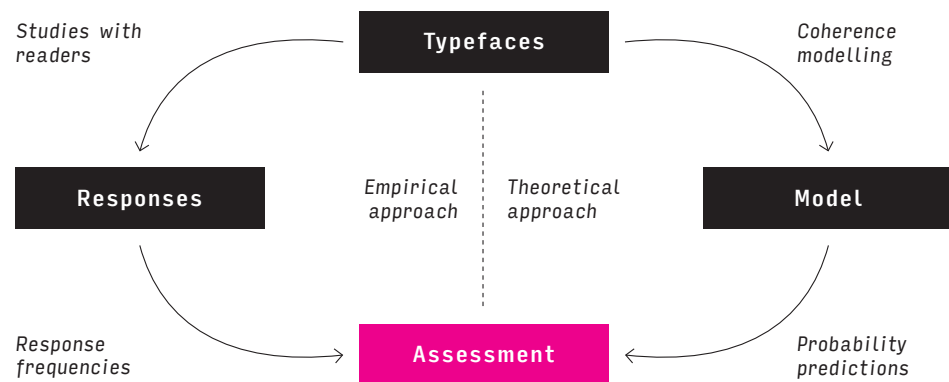


FIGURE 1.6: The design of the research consists of studies with participants (empirical approach) and modelling character similarity and coherence (theoretical approach). Results from both of these components are compared in the final assessment of the model.

1.3. Impact

One of the aims of this research is to develop a theory for a systematic analysis of coherence. A theory that would be related to perception as well as typeface design practice and could provide a basis for a subsequent formal perspective of typeface design. A perspective that could be used alongside the already well-established historical and technological views. More generally, the study of coherence can contribute to any field dealing with the question “What makes an ‘a’ an ‘a’?”. The font-tuning studies mentioned above (Sanocki, 1987; Sanocki, 1988; Sanocki, 1992; Gauthier et al., 2006; Walker, 2008; Sanocki & Dyson, 2012) as well as other studies (Palmer, 1999) show that the recognition of characters is influenced by context, both visual and linguistic (see figure 1.7). In this respect, the study of coherence can be seen as a study of the visual context created by typefaces and their designers. Although not providing the complete explanation of the processes involved in character recognition, better understanding of coherence could help to address the question about the essence of characters which is an important topic in epigraphy, linguistics (graphemics), psychology (letter-recognition and readability research), and cognitive science.

THE CAT

FIGURE 1.7: Stephen Palmer (1999, p. 429) shows that the identification of characters is influenced by context. English readers should not have a problem reading the two words above as “THE CAT”. Note, however, that during the reading process, a single shape has been identified as two different characters depending on its context. H in “THE” and A in “CAT” both use the same shape. This shows that lexical knowledge can be an important factor in character identification, among others.

One of the challenges in readability research is the verification or rejection of the aforementioned question of whether coherent typefaces are easier to read. A theory of coherence that would correspond to typeface-design practice and allow measurement of coherence should help with articulating the question in a formal and testable way.

From a typeface-design perspective the theory could feed into font production processes, typeface design education, quality assessment of fonts, as well as multi-script typeface design. A few directions for possible contributions are suggested in the following paragraphs.

Greater formality for discussing coherence supported by empirical evidence will help to make communication of visual relationships between characters clearer. This will be beneficial in typeface design education, as

well as in practice, especially when dealing with large projects that require a coordinated design effort. The methodology described in chapters 5 and 6 could also help in situations where design conventions and processes for particular scripts are currently not well documented.

In his article *Lessons learned from Metafont*, Donald Knuth (1985) states that Metafont's "ability to capture the designer's intentions rather than just the drawings that result from these intentions [...] will prove to be much more important than anything else". He seems to be referring to intentions visualized in character drawings, but not necessarily captured in a formal way. While it is questionable whether any computer system can ever capture the sheer complexity of any designer's intentions, Knuth's observation that typeface design is more than a set of independent drawings is acute. A tool that ought to assist designers during their creative process might, and perhaps should, go beyond mere assistance with construction of individual drawings. Coherence that spans across different character shapes definitely belongs among designers' intentions which are not usually captured in font production tools.

Understanding coherence as it is perceived by readers will help designers to better distinguish their design-specific ideas from generally accepted and expected similarity relationships. This will lead to a greater effectiveness in their creative process.

Contemporary font-development software might implement the presented approaches on two levels. First, tools could provide real-time auditing and suggestions by comparing the work in progress with predetermined script conventions. Second, tools could use these conventions in an implicit way to generate new character shapes based on the ones that have already been designed. A reasonable implementation of such an approach is up for discussion. On the one hand, software tools could become more effective and remove or at least alleviate some of the tedious and repetitive aspects of typeface design work. On the other hand, each typeface design tool must consider design freedom and intrinsic variety of script conventions. Insensitive application of the theory might lead to overly restrictive tools that make interrogating conventions more challenging if not impossible.

Independence from technology is an important attribute of the theory as it allows for discussion of typefaces and their appearance across different platforms. Although contemporary font formats promise to be transferrable across various reproduction environments, little control over their appearance in these environments results in incoherencies that may affect reading experience. This issue of translatability of typeface design has been discussed in detail by Richard Southall (1986, 1991, 2005). The ability to encode coherence among characters may be the first step for typeface appearance to become

truly cross-platform and consequently ready for any future technologies.² Further, it will improve quality assessment by allowing comparison of some of the designer's intentions with the reproduced results.

Better understanding of coherence is also of great importance in multi-script typeface design, for which one of the key questions is a stylistic harmonisation of two or more scripts. Formalizing principles of coherence in scripts individually and in a coordinated way should establish the conditions for addressing the script harmonization challenge, and improve methodology and discourse within multilingual and multi-script typography.

1.4. A few notes about the scope and style of this thesis

When writing this thesis, the main challenge lay in the fact that it is cross-disciplinary. However, this also constitutes what I hope will end up being its important contribution. Chapter 2 sets out terminology to make the work relatable to linguists, psychologists, and typographers alike. Chapter 3 clarifies the basics of the scripts studied both from the linguistic and typographic point of view. Chapter 4 contains an overview of the common typeface design approaches and tools for conceptualizing character shapes. These shorter chapters set the context while the central focus of the thesis is in the following two chapters. Chapter 5 describes the testing methodology and statistical analysis used for the studies with readers. Finally, chapter 6 uses data analysis with a little bit of machine learning to explain the cognitive processes involved in judgements of character similarity. This summary promises an exciting mixture of topics and perspectives. However, this mixture also makes it difficult to write a text that is both relevant and digestible to typographers, typeface designers, linguists, psychologists, and data scientists. I have made an effort to write a text that is as readable as possible in parts where approaches from potentially unfamiliar fields are presented, but that is not overly simplistic in parts where familiar approaches are discussed. More technical considerations that are used to support the main argument and are not essential for the theoretical narrative are extracted into individual sections, footnotes, or appendices, so those who are not keen on reading them can skip them easily and not lose the thread of the central argument.

2. For example, rasterization of contemporary fonts on low-resolution screens sometimes results in an incoherent treatment of features that should remain coherent, e.g. the thickness of strokes or serifs varies across different characters from the same typeface. In contemporary font formats based on contour descriptions this issue is partially solved by the introduction of an additional layer of hinting instructions. Even though they solve the rasterisation problem sufficiently, the hinting instructions are still an incomplete description of typeface coherence. Note that rasterization does not seem to be an issue in Metafont which deals with the issue in a more comprehensive way. Nonetheless, the various issues of rasterization are not the subject of this research.

1.5. Note on notation and figures

Throughout this thesis, specific character shapes the readers should look at as if they were illustrations are marked in text using a grey background. Note that **a, b, c** refers to a set of shapes (a triplet in this case) that should be considered together while **a, b, c** refers to three shapes that should be considered individually. When referring to general character categories or their associated sounds, single quotation marks are used in a similar fashion: 'a, b, c' for a set and 'a', 'b', 'c' for three individual categories. In case there is a need for a reference to a character in a figure, single quotes are used. The brackets are used for transliteration or pronunciation.

Unless stated otherwise, images from other publications are reproduced at 100 % scale.

2. Theory and definitions

The fields that deal with subjects related to this research – fields such as cognitive psychology, linguistics, computer science, epigraphy, and typography – use varied and often incompatible terminology. In order to allow discussion of character similarity, typeface coherence, script conventions, and other ideas presented here, this chapter puts forward terminology that is based on existing notions within the aforementioned fields, but redefined to provide a clear perspective that is based on the visual aspects and accounts for multi-script and multilingual considerations. Corresponding terms from other fields are mentioned where appropriate.

2.1. Character

The term *character* is problematic as it is not always clear from its definitions whether it refers to a visual or non-visual entity. *The Oxford English dictionary* (2018) entry for the term is extensive, but effectively it defines character either as “a member of a set of symbols used in writing or printing to represent linguistic elements, as individual speech sounds, syllables, or words; any of the simple elements of a written language, as a letter of an alphabet, or an ideogram.” or, in computing, also as “a representation of such a symbol by means of a small number of bits, holes in punched tape, etc.”. *The Unicode standard* (2014, p. 6) states the character identified by a Unicode code point is an abstract entity (e.g. Devanagari digit two) and the mark made on screen or paper, called a glyph, is a visual representation of the character. *The Unicode Standard* definition and the *Oxford English dictionary* definition for computing are non-visual. On the other hand, in typography or epigraphy, characters are usually visual representations.¹ This work adheres to this meaning as well.

Another question is whether a character is a general or specific entity. This is resolved here by stating that these are essentially two different modes of use of the same term. In a wider sense, a character is general and refers to a category of symbols representing the same information, typically a linguistic unit, e.g. all letters ‘b’ or any letter ‘b’. In the most narrow sense, it refers to a specific symbol, a material object, e.g. the letter **b** that you have just read. To distinguish between these two modes of use when necessary, two compound terms are used: the general *character category* and the specific *character shape* (see figure 2.1). The set of all character shapes from the same character category is called a *character category domain*.

1. To be precise and not to discriminate against blind people, one should in fact speak of visual or tactile representation in typography or epigraphy. However, this work focuses only on the visual perception of characters.

It is worth noting that character shapes from the same category can look very different depending on their styling (see figure 2.2), colour, scale, distance, geometric and other manipulations (e.g. those provided by DTP software), but also depending on the reproduction technology used. Various kinds of screens, rasterizers, and printers each have different effects which show in the resulting shapes.




FIGURE 2.1: Two modes of using the term character. General character categories (illustrated as points of the cones) contain specific character shapes. The character category domain (bases of the cones) comprises all character shapes that can be recognized as belonging to the character category.

Importantly, a character always represents a conventional linguistic unit. Based on the structural level of the associated linguistic unit, one can distinguish several kinds of characters (Coulmas, 1989). There are three kinds that are used in this research:

- i. a *letter* is a character representing one or more of the sounds used in speech, e.g. the Latin-script letter **b**,
- ii. a *syllabic character* (or *syllable*), is a character representing a syllabic sound, e.g. the Devanagari syllables **ब** [Ba] or **ड्ड** [DdDda],
- iii. an *orbital sign* is a character that cannot be used on its own, but can be combined with other characters to modify their meaning, e.g. the caron **ˇ** accent as used in letters such as **ř** or the Devanagari vowel sign **ी** [-ii] as used in combinations such as **पी** [Pii].

This list is not meant to be complete. It leaves aside logograms or ideograms such as those used in Chinese and Japanese scripts, for example, because these scripts are not discussed in this work.

In a multilingual environment, the situation is more complex as the same character can represent different linguistic units depending on a language, e.g. the character shape **c** represents a different sound in English compared to Czech. Nevertheless, this seems to be a linguistic issue as the identity of a character remains the same: **c** would be identified as the same symbol by both

English or Czech readers no matter what language context it were used in. In other words, a character is a token associated with a linguistic unit, but this association is merely conventional (see also Saussure, 1972). The linguistic unit itself does not prescribe the appearance of its representative shape nor vice versa, the script conventions do. Should the **c** go missing from the repertoire of a favourite font, another symbol could be used in its stead if this was agreed upon (conventionalized).

It is also important to note that an uppercase (or capital) letter and a lowercase (or small) letter representing the same sound are considered two distinct character categories in this research. In other words, the linguistic unit they represent is different. For example, in the Latin script, character shapes **A** and **a** belong to two different categories, while **á** and **à** belong to the same category.

There are similar terms in other fields. The term character category roughly corresponds to the linguistic term *grapheme* when this is defined as the smallest functional unit of writing (Coulmas, 1996, p. 174). On the other hand, cognitive psychology often uses the term *abstract letter identity* (see for example: Kinoshita & Kaplan, 2008) which seems to refer to all letters representing a particular sound regardless of case, i.e. it could include multiple character categories. For example, **A** and **a** would have the same abstract letter identity.

2.2. Character concept

Readers that are not very familiar with typeface design and typography often assume that there is a single shape-like idea prescribing the appearance of all character shapes from a particular category. However, after careful consideration, a single invariant prototype based on shapes seems untenable.

There are several problems that make the task of describing all shapes in a character category really difficult. Firstly, character categories, although illustrated here with simple diagrams with well-defined borders between them, are ever-evolving social constructs. Thus, they also depend on perceptions of individual readers. What might pass as an 'a' for some, others might not perceive as one. Secondly, designers keep on producing new typefaces and letterings. New radical designs constantly push the boundaries of what are acceptable members of character categories. Thirdly, characters exist in the context of other characters. Sometimes knowing how the other characters look may help identify previously unseen characters, especially in less formal typefaces where characters might be easy to confuse. And lastly, the linguistic context has an effect on character identification as well. What is an 'H' in one context, can be an 'A' in a different context (see figure 1.7).² Thus, describing all the shapes in a character category is a task that

2. Note that character categories can also overlap when the same shape can be interpreted as belonging to two different character categories.

would involve not only describing the social construct, it would also require a working model of human perception and creativity in the context of character design, including all possible production processes. Clearly, this would be a difficult task.

Douglas Hofstadter (1982) challenged an idea that all character shapes belonging to a particular category could be described with a parametric shape-based prototype on the grounds of mathematical logic. Illustrating the difficulty of parametric approach with many variants of an 'a' (see figure 2.2), he suggested an interconnected web of general criteria that make up the idea of 'a' itself (Hofstadter, 1982). The exploration of the idea of 'a' and the automated creation of a coherent style were some of the objectives of the Letter Spirit project lead by Hofstadter within the Fluid Analogies Research Group in the 1990s and early 2000s (Hofstadter & McGraw, 1998; Rehling, 2001).

While it seems unfeasible to describe the whole domain of all shapes from a particular character category, it is possible to describe at least some subsets of the domain in order to make general comments about the character shapes included. Subsets of character shapes from the category domain can be described using *character concepts*. A character concept is yet another mode of using the term character. Essentially, it is any generalized description of character shapes. A character concept is not as general as a character category and not as specific as a character shape. There is no assumption being made regarding an approach taken to describe the shapes other than it has to focus on the visual aspects. It can be more or less general. The more general the concept, the larger the set of character shapes it describes, and so the larger the *character concept domain* (see figure 2.3).

Some character concepts can be used in the production of character shapes, e.g. in fonts which determine the shapes of characters in the context of reproduction technologies. However, other character concepts might not be suited for production, especially those that are very general. That being said, character concepts can always be used to judge whether a character shape belongs to their domain, that is, whether its shape complies with the concept's definition.




FIGURE 2.2: The variety of shapes among 70 different letters 'a' shows how challenging it is to find a unifying description for all possible shapes of 'a' (from: Hofstadter and McGraw, 1998, p. 413).

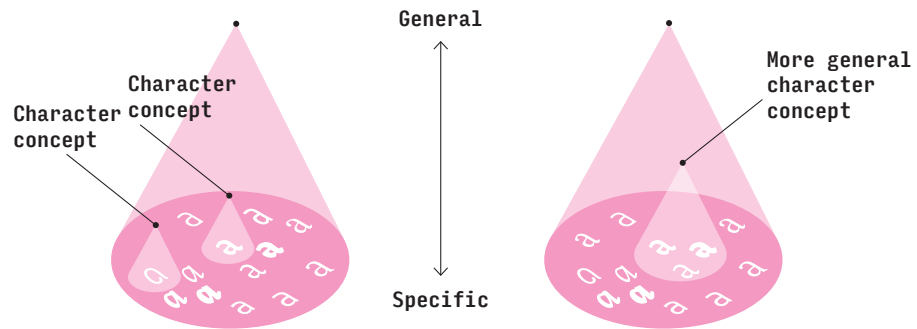


FIGURE 2.3: The domains of all potential character shapes (bases of the large cones) for particular character categories (points of the large cones) are compartmentalised by character concepts (points of the small cones) into smaller character concept domains (bases of the smaller cones).

For example, the character shape **a** could be conceptualized as “having an arch on the top that is smoothly connected to a vertical stem on its right side with a small loop tucked underneath the arch on the left side of the stem ...” either verbally or in some kind of formal system. Clearly, many other shapes of an ‘a’, would satisfy this fairly general description and thus fit within the domain of this concept.

On the other hand, an ‘a’ could be conceptualized by means of a much more precise geometric description of its contours, e.g. bézier curves in an OpenType font (Microsoft, 2018), which would significantly limit the number of all potential shapes of ‘a’ fitting this concept (see the ‘a’ on the left in figure 2.4).

It is worth noting that although geometric character concepts may seem very precise, they can still realise a set of different character shapes depending on the reproduction technology used. This has been thoroughly discussed by Richard Southall (2005) and others. Some coarse printers, for example those used in cash registers to print receipts, may produce very different results from high-resolution screens based on the same concept.

The geometric descriptions can be more general too: the concept of a letter ‘a’ in a variable OpenType font (Microsoft, 2018) that allows variation spanning from regular, to bold, and to black weight variant would be implemented differently from the example in the previous paragraph. The contours of the letter would be described using bézier curves and a set of values that specify how these curves change between the weight variants (see the three ‘a’s on the right in figure 2.4). The set of realised character shapes contains a multitude of weight variants of ‘a’ between regular, bold and black and their interpretations in various reproduction technologies. Thus, this character concept is more general.

Note that the more general character concepts can theoretically include other character concepts. In the examples above, all the character shapes of an ‘a’ realised using the single-style OpenType font can also be realised using the character concept for an ‘a’ from the Variable OpenType font (compare the concepts on the left with the concept on the right in figure 2.3 and both with figure 2.4).

Various ways of conceptualizing characters used in typeface design and font production will be discussed in chapter 4: *Approaches to character conceptualization*.

The idea of character concept is related to Debra Adams and Richard Southall's (1989) term "character appearance specification", Southall's (1991) terms "model" and "pattern", and McGraw's (1995, pp. 6–7) terms "letter concept, letter conceptualization, and letter plan". All of these are character concepts with varying levels of generalization.




FIGURE 2.4: Single-style OpenType fonts (the single 'a' on the left) represent character concepts as a set of bézier curves. Variable OpenType fonts (the three 'a's on the right) represent character concepts as a set of bézier curves for one character with associated values that specify how these curves change between character variants (illustrated by the arrows), in this case changing the weight of the characters. Only some of the control points (marked in magenta) defining bézier curves are shown and the contours are dotted to illustrate that only the control points are included in the concept, not all the points on the curves. Only some of the arrows indicating the variability are shown here for the sake of simplicity.

The model proposed in chapter 6 uses simple character concepts based on expert knowledge in an attempt to describe criteria readers might use when judging similarity of character shapes. Clarifying how these concepts relate to readers' similarity judgements will help answer the third research question: whether the coherence perceived by readers differs from the coherence intended by designers (see chapter 1). Understanding the difference in perception of coherence between designers and readers is an important step towards a better understanding of perception of typefaces.

2.3. Writing and script

The term *writing* in this work includes any use of characters to convey a message, in print, on screen, or in handwriting. For writing to be deciphered, one has to know the linguistic units conventionally associated with characters. In the words of Gerrit Noordzij (1973, p. 80): “Writing is a convention. [...] Unconventional writing is no writing at all.”

Writing system and *script* are two terms that are often used interchangeably although they can have distinct meanings. This work adheres to the definitions provided by Florian Coulmas (1989). He uses the term *writing system* to differentiate systems depicting linguistic units of different structural levels, e.g. alphabetic writing system, syllabic writing system, etc. Following his definition, the term *script* refers to a visual instantiation of a writing system, e.g. Latin script, Devanagari, etc. Further, he notes that the term *script* is often used as if scripts were inherently related to a given language and that this is only partly true. There are cases where one script is used by multiple languages, e.g. Devanagari is used for Hindi, Marathi, Nepali, and other languages. Admittedly, there are cases where there is just one language using a given script, e.g. Greek, or cases where one language uses multiple scripts, e.g. Serbian uses Latin as well as Cyrillic scripts. In addition, the rules for applying a script (its *orthography*) to a particular language may vary as well.

For the purpose of this work, and in line with Coulmas’s notions, a script can be defined as a set of character categories used for writing, together with the general system of their use, i.e. orthography (see figure 2.5).³ As the focus is on the visual, the scripts’ orthographies will be discussed only briefly in chapter 3: *Overview of the scripts studied*.

3. The definition used in this work interprets a script as a contrast set. A contrast set is a set of categories that are mutually exclusive, i.e. contrasted, (Watson, 2005). Contrast sets are commonly used to describe various domains of human interest, e.g. temperature scales, colour palettes, countries, alphabets, and therefore should be a straightforward concept to work with. Using the concept of a contrast set will also allow for a clearer discussion about how individual character categories are formed and how they relate to each other.




FIGURE 2.5: An example of a script represented as a set of character categories. The Devanagari script is represented here with only a few character categories and shapes. The orthography is not illustrated.

2.4. Typeface

According to Adams and Southall (1989, p. 213), “a typeface is a set of appearance specifications for the objects that realise the characters of a script.” Note that the term character as it is used in their article refers to what is called a character category in this work. Defining a typeface this way is practical as it clearly links characters to typefaces and to scripts. This allows for descriptions of relationships between characters at different levels of generalization.

Adapted to the already established terminology and rephrased to work within the context of multi-script typography, Adams and Southall’s definition would read as: A typeface is a set of character concepts, each of them specifying character shapes for one of the character categories from one or more scripts (see figure 2.6).⁴ Notably, a typeface defined this way can contain character concepts from multiple scripts. For example, the typeface Arial contains character concepts used for Latin and Cyrillic scripts.

Character concepts for a typeface are typically implemented in a *font*. The approach to conceptualisation is usually shared for all character concepts implemented by a font and is determined by the *font format*.

Character concepts included in typefaces and fonts are the results of typeface designers’ work. It is an inherent part of typeface designer’s job to make sure the character concepts determine character shapes that can be recognized as belonging to the corresponding character category. Their failure at this task leads to less readable fonts which may lead to readers misreading the intended message.

4. It is worth noting that some designers (e.g. Bigelow, 1985) consider typefaces as coherent by definition. With such a definition, to be able to say precisely what is and what is not a typeface would require a definition of coherence and the criteria used to assess the presence or level of coherence in a typeface. In order to avoid any assumption regarding coherence, this research uses a definition where coherence is not considered a necessary quality of a typeface.

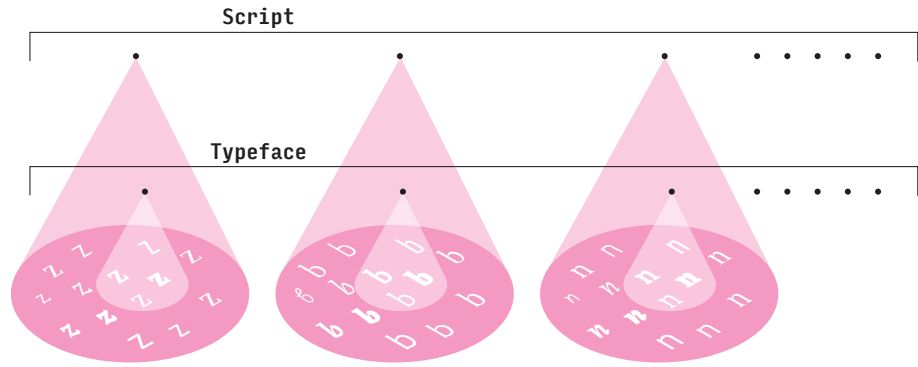


FIGURE 2.6: Relationships between a script consisting of character categories (points of the large cones), a typeface consisting of character concepts (points of the small cones), and character shapes (within the bases of the cones).

To avoid ambiguity, it is worth mentioning that the term *typeface* is commonly used in a broad sense by designers. The scope of this use has been thoroughly discussed by Robin Kinross (2002) without a definitive resolution. Fred Smeijers’s (1996, p. 16) definition of a typeface as “a set of types over a whole range of sizes” seems to be the most generally accepted. According to him, the idea of a typeface emerged gradually around the 16th century when individual fonts for specific sizes began to be joined by a consistent, unifying notion of their forms. Later, variants, such as italic, small capitals, bold, light, and more, were included within the notion of a typeface as well. Thus, the term came to be used as a user-oriented umbrella term to suggest design coordination or compatibility of its constituents. To avoid confusion, the term *type family* is used for this broader meaning of the term (see figure 2.7).

Arek Regular	<i>Arek Italic</i>
Arek Semibold	<i>Arek Semibold Italic</i>
Arek Bold	<i>Arek Bold Italic</i>
Arek Extrabold	<i>Arek Extrabold Italic</i>

FIGURE 2.7: The type family Arek contains multiple variants, e.g. uprights (left) and italics (right) in the following weights (top to bottom): Regular, Semibold, Bold, Extrabold.

2.5. Similarity and coherence

Where similarity refers to relationships between individual objects, coherence refers to an overall quality. In this work *coherence* is viewed simply as a structure of individual similarity relationships between characters. Thus there can be a coherence of a group of character shapes, coherence of character concepts (i.e. *typeface coherence*), and coherence of character categories. The conventional coherence of character categories from the same script is called *script conventions*.

Defining characters as visual entities with different modes means that conclusions can be made about character concepts within a typeface based on the evidence about material character shapes. And analogically, conclusions about general character categories within scripts can be based on evidence about the character concepts or character shapes. Thus, the similarities between character shapes can be aggregated and taken as evidence of inter-concept and inter-category similarity relationships and coherence (see figure 2.8). Note that none of this would be possible if characters or character categories were defined as abstract entities.

The key to understanding script conventions and typeface coherence is character similarity which is studied empirically in chapter 5 and theoretically in chapter 6.




FIGURE 2.8: Similarity relationships between character shapes constitute the basis for similarity relationships between character concepts in typefaces and conventional similarity relationships between character categories from a script.

2.6. Style

The term *style* is commonly used with various meanings. It may be used to suggest something exceptional (“This typeface has style!”); it may refer to a particular period of time (“It has been designed in the art deco style.”); it may refer to an artist’s personal voice (“He has a peculiar style.”); or it may refer to a particular variant of a typeface or a single-style font from a type family (“Let’s use the bold-style font for the headline.”). In addition, it may be understood as a non-essential part of a shape’s appearance, particularly in visual communication and web design where styles are often defined independently of the shapes or content they are assigned to (“Make sure the footnotes and sidenotes have the same style applied to them, including the colour and rounded section border.”). The many different usages of the term *style* make it difficult to use in academic writing. Nonetheless, it is an important term in design practice and it is useful to relate the present study of coherence to it.

For visual arts and for art history in particular, style is one of the central issues and there are several key texts that discuss it. In particular, Ernst Gombrich (2009) and Meyer Schapiro (1994) have achieved sufficiently general definitions which permit their transfer to typeface design.

Gombrich (2009, p. 129) defines style as: “a distinctive, and therefore recognisable, way in which an act is performed or an artifact made or ought to be performed and made”. The stress here is on recognizability and distinctiveness rather than on similarity to other acts or artefacts. Gombrich develops this further by quoting Stephen Ullman (1957, p. 6): “The pivot of the whole theory of expressiveness is the concept of choice. There can be no question of style unless the speaker or writer has the possibility of choosing between alternative forms of expression.”⁵ In this view, an object can have a style only when it is clearly distinct from others it is compared to and when there is a variety of alternative forms—this is referred to as a *stylistic space* in this work.

In order to compare objects or forms to say whether they are similar or distinct, there has to be a set of criteria that allows for the comparison—a system that allows describing different objects in a compatible way in order to compare them. For example, one can imagine comparing a chair and table by saying they share the same upholstery and wood carving. In this case, the different pieces of furniture have been described as “objects with upholstery and carving” and that permits further discussion. It might

5. The word ‘choice’ may seem to suggest an author’s ultimate understanding of the final object and control over its interpretation. It was already discussed in the introduction that this is not always the case. There is a difference between the intention of the designer and the intention of the reader.

seem that nothing in Gombrich’s definition suggests that style can link objects together, but the need for comparison implies some form of linking.

Schapiro (1994, p. 51), on the other hand, puts the requirement for similarity in a more straightforward way: “By style is usually meant the constant form—and sometimes the constant elements, qualities, and expression—in the art of an individual or a group.” Investigating style, according to Schapiro (1994, p. 59), is “often a search for hidden correspondences explained by an organizing principle that determines both the character of the parts and the patterning of the whole.” And in the words of typeface designers Erik van Blokland and Just van Rossum (2004), “style in the broader sense could be the art of finding common threads in seemingly unrelated things and defining groups for them.”

Thus style is defined in this work as a set of distinctive visual attributes of forms which permits their grouping with other forms. Style in typeface design is often described by means of typographic attributes (see section 4.5.1) such as the differences between thin and thick parts of the strokes (the contrast), the weight of the thick parts, height, etc. Nonetheless, these attributes are not applied equally to all characters from the same typeface (see figure 2.9). Different character categories make different allowances for the implementation of the global attributes and style as described this way. In order to understand style better, one ought to study individual characters and the attributes they can carry. In this respect, coherence can be seen as a structure describing the ways the attributes of a particular style are implemented in different characters.

Approached from a different point of view, typefaces do not have to be coherent, but their coherence is often imposed on them by a designer. And one of the motivations to do so, as discussed already in the introduction, is the appeal of a uniform style. In order to produce a uniform style, designers need to distribute the attributes of style consistently across the different character concepts. Coherence, in this sense, is the result of the work of a designer: a coordination of character concepts in order to produce coherent character shapes.



FIGURE 2.9: Global attributes are not applied equally to all characters from the same typeface. The tapering of strokes is present in **i**, **n**, **x**, but not in **o**. The difference between the thin and thick strokes is most apparent in **o**, **x**, less so in **n**, and not present in **i**. The example is in the typeface Optima.