Preface

The last ten years have witnessed the fact that geometry, topology, and algorithms form a potent mix of disciplines with many applications inside and outside academia. We aim at bringing these developments to a larger audience. This book has been written to be taught, and it is based on notes developed during courses delivered at Duke University and at the Berlin Mathematical School, primarily to students of computer science and mathematics. The organization into chapters, sections, and exercises reflects the teaching style we practice. Each chapter develops a major topic and provides material for about two weeks. The chapters are divided into sections, each a lecture of one and a quarter hours. An interesting challenge is the mixed background of the audience. How do we teach topology to students with a limited background in mathematics, and how do we convey algorithms to students with a limited background in computer science? Assuming no prior knowledge and appealing to the intelligence of the listener are good first steps. Motivating the material by relating it to situations in different walks of life is helpful in building up intuition that can cut through otherwise necessary formalism. Exposing central ideas with simple means helps, and so does minimizing the necessary amount of technical detail.

The material in this book is a combination of topics in geometry, topology, and algorithms. Far from getting diluted, we find that the fields benefit from each other. Geometry gives a concrete face to topological structures, and algorithms offer a means to construct them at a level of complexity that passes the threshold necessary for practical applications. As always, algorithms have to be fast because time is the one fundamental resource humankind has not yet learned to manipulate for its selfish purposes. Beyond these obvious relationships, there is a symbiotic affinity between algorithms and the algebra used to capture topological information. It is telling that both fields trace their names back to the writing of the same Persian mathematician, al-Khwarizmi, working in Baghdad during the ninth century after Christ. Besides living in the triangle spanned by geometry, topology, and algorithms, we find it useful to contemplate the place of the material in the tension between extremes such as local vs. global, discrete vs. continuous, abstract vs. concrete, and intrinsic vs. extrinsic. Global insights are often obtained by a meaningful integration of local information. This is how we proceed in many fields, taking on bigger challenges after mastering the small ones. But small things are big from up close, and big things are small from afar. Indeed, the question of scale lurking behind this thought is the

driving force for much of the development described in this book. The dichotomy between discrete and continuous structures is driven by opposing goals: machine computation and human understanding. The tension between the abstract and the concrete as well as between the intrinsic and the extrinsic has everything to do with the human approach to knowledge. An example close to home is the step from geometry to topology in which we remove the burdens of size to focus on the phenomenon of connectivity. The more abstract the context the more general the insight. Now, generality is good, but it is not a substitute for the concrete steps that have to be taken to build bridges to applications. Zooming in and out of generality leads to unifying viewpoints and suggests meaningful integrations where they exist.

While these thoughts have certainly influenced us in the selection of the material and in its presentation, there is a long way to the concrete instantiation we call this book. It consists of three parts and nine chapters. Part A is a gentle introduction to topological thought. Discussing graphs in Chapter I, surfaces in Chapter II, and complexes in Chapter III, we gradually build up topological sophistication, always in combination with geometric and algorithmic ideas. Part B presents classical material from topology. We focus on what we deem useful and efficiently computable. The material on homology in Chapter IV and duality in Chapter V is exclusively algebraic. In the discussion of Morse theory in Chapter VI, we build a bridge to differential concepts in topology. Part C is novel and the reason for why we wrote this book. The main new concept is persistence, introduced in Chapter VII, and its stability, discussed in Chapter VIII. Finally, we discuss applications in Chapter IX.

<div align="right">

Herbert Edelsbrunner and John L. Harer
Durham, North Carolina, 2009

</div>