

Geometric Abstraction for Effective Visualization and Modeling

DISSERTATION

zur Erlangung des akademischen Grades

Doktor der Technischen Wissenschaften

eingereicht von

Dipl.-Ing. Haichao Miao, BSc

Matrikelnummer 00726810

an der Fakultät für Informatik
der Technischen Universität Wien

Betreuung: Assoc. Prof. Dipl.-Ing. Dr.techn. Ivan Viola
Zweitbetreuung: Mag. Dr. rer.nat. Ivan Barišić, Dipl.-Inf. Dr.-Ing. Tobias Isenberg
& Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Eduard Gröller

Diese Dissertation haben begutachtet:

Prof. Dr. Helwig Hauser

Prof. Dr. Ingrid Hotz

Wien, 5. August 2019

Haichao Miao

Geometric Abstraction for Effective Visualization and Modeling

DISSERTATION

submitted in partial fulfillment of the requirements for the degree of

Doktor der Technischen Wissenschaften

by

Dipl.-Ing. Haichao Miao, BSc

Registration Number 00726810

to the Faculty of Informatics

at the TU Wien

Advisor: Assoc. Prof. Dipl.-Ing. Dr.techn. Ivan Viola

Second advisor: Mag. Dr. rer.nat. Ivan Barišić, Dipl.-Inf. Dr.-Ing. Tobias Isenberg

& Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Eduard Gröller

The dissertation has been reviewed by:

Prof. Dr. Helwig Hauser

Prof. Dr. Ingrid Hotz

Vienna, 5th August, 2019

Haichao Miao

Erklärung zur Verfassung der Arbeit

Dipl.-Ing. Haichao Miao, BSc
Huttengasse 18/3/13, 1160 Wien, Austria

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

Wien, 5. August 2019

Haichao Miao

Acknowledgements

After almost four years, an important chapter of my life will close with this thesis. It has been a wonderfully challenging journey, with many people that accompanied me along this way. All of them contributed to the development of this thesis in one way or another and I want to express my sincere gratitude for them.

First and foremost, I would like to express my deep gratitude to my supervisor Ivan Viola, whose high demands are only matched by his great support. Ivan, thank you for providing me with the independence I wanted and giving me the guidance I needed. You had lots of good ideas and ensured a smooth and thriving environment for my research!

I am indebted to my co-supervisor Meister Eduard Gröller. Our countless meetings have been a pleasure and I benefited so much from your sound advice and encouragement. I was very fortunate to work in the unique environment you established. I'd like to thank all the current and former members from this environment, the VisGroup: I thank Gabriel, who introduced me to the visualization field and supported me at the beginning of my PhD with lots of valuable and appreciated advice. Yun (your playful curiosity will always remain an inspiration to me), MedVis-pro Renata (the helpful and cheerful soul), Alexandr (forces you to smile :), rockin' Manu (and your endless knowledge on evaluation), Peter (let's jam again sometime!), Tobi (best ski instructor), Bara (I hope to visit you in Nepal), Sarkis (soad4ever!!), Johannes (YourHighness, Johnny, Jackson...), Johanna, Alexey, Viktor, and Nicholas. Special thanks go to my buddy, David. I will miss our daily deep and not-so-deep discussions.

I wish to thank my co-supervisor Tobias Isenberg at Inria, Paris. I was inspired by your way of formulating and shaping ideas into papers. I learned so much working together with you and enjoyed the support and guidance that you dedicated to me!

I would like to thank the examination committee for their time and efforts, especially the two International Experts, Ingrid Hotz and Helwig Hauser. I appreciate the time you took to review this work and your insightful feedback. I thank Michael Krone, the National Expert, for his valuable input he provided early on.

I was extremely lucky to have collaborated with so many supportive researchers. Without them, I would not have been able to develop the ideas in the papers of this thesis. I thank the doctors at the King's College for providing the necessary input and feedback, Christian Nasel from Karl-Landsteiner Privatuniversität, who supported me with data

of stroke patients, ideas and exciting discussions in many long meetings, and Bernhard Kainz from Imperial College, who provided the data, segmentations and the appreciated input for the PlacentaMaps project.

During my PhD, I was also part of the Molecular Diagnostics team at the AIT, headed by Ivan Barišić. I was so lucky to be co-supervised by you, Ivan. You have always been supportive, motivating and found the time to discuss ideas. I enjoyed your positive attitude, your professional as well as personal advice and I'm sure they all will have an lasting impact on me. Thank you! I would also like to thank my colleagues Elisa, Michi (I was never so wrong with my first impression :), the always good-humored Ari, super-knowledgeable Noa, Timo (thanks for the three strawberries :P), Rick, Tadija, Stephan, Yasi, and Regi. I cannot thank Elisa De Llano enough, who I collaborated with daily for a major part of my PhD. It was a pleasure working with and learning from you.

I was a visiting researcher at Harvard University in the Visual Computing Group, where I met a lot of great people. Big thanks to Hanspeter Pfister and Johanna Beyer for giving me this great opportunity, and to Ivan and also Ivan, who supported me in this endeavor. It was an amazing time rich on unique experiences.

I'd also like to thank the secretaries Anita and Max for the organizational help, and I thank the technicians Stephan, Andi, and Simone, who always ensured a smooth operation in the background.

我感谢我的父母无条件的爱与支持。我更感激我的父母提供我这样的机会和经验，进而造就了现在的我。 Ich danke meiner kleinen Schwester Jen, die mich früh gelehrt hat verantwortungsvoll zu sein. Danke für die vielen Male, in denen du als 姑姑 eingesprungen bist.

Zuletzt möchte ich mich bei meiner Partnerin Conny für ihre uneingeschränkte und liebevolle Unterstützung bedanken. Du hast von Anfang an an mich geglaubt und mir gezeigt, dass alles doch relativ ist. Ich widme diese Arbeit meinem Sohn Peter, der meine größte Antriebskraft ist und mir stets vor Augen hält, dass alle Anstrengungen sich lohnen.

Kurzfassung

In der vorliegenden kumulativen Dissertation wird geometrische Abstraktion als eine Strategie zur Generierung eines integrierten Visualisierungssystems für räumliche wissenschaftliche Daten beschrieben. Die hier behandelte Herangehensweise abstrahiert räumliche Daten auf zwei Arten. Diese Methode erschafft eine Vielzahl an Repräsentationen. Entlang der Achse für Räumlichkeit werden räumliche Details graduell entfernt, während die Merkmale entlang der Achse für visuelle Details zunehmend aggregiert und durch unterschiedliche visuelle Objekte repräsentiert werden. Diese Darstellungen werden in weiterer Folge in einen konzeptionellen Abstraktionsraum integriert, welcher es Benutzern erlaubt, mittels effizienter Steuerung der Repräsentation das Abstraktionslevel an eine gewünschte Aufgabe anzupassen. Um es dem Experten zu ermöglichen, Übereinstimmung zwischen diesen Repräsentationen herzustellen, werden animierte Übergänge eingesetzt. Schlussendlich speichert der Abstraktionsraum Benutzerinteraktionen und bietet visuelle Indikatoren an. Damit wird der Experte in Richtung interessanter Repräsentationen für eine bestimmte Aufgabe und Datensatz verwiesen. Mentale Modelle der Experten spielen eine wichtige Rolle im Verstehensprozess abstrakter Repräsentationen. Sie werden daher beim Design des Visualisierungssystems berücksichtigt, um kognitive Beanspruchung des Benutzers gering zu halten. Der hier vorgestellte Ansatz wird anhand der beiden unterschiedlichen Felder der Plazentaforschung und des *in silico* Designs von DNA Nanostrukturen demonstriert. Beide Forschungsfelder nutzen geometrische Abstraktion zur effektiven visuellen Untersuchung sowie für spätere Modellierungsaufgaben. Das *Adenita* Toolkit, eine Software zum Entwerfen neuer DNA Nanostrukturen, implementiert die hier vorgestellten Visualisierungskonzepte. Dieses Toolkit, zusammen mit den vorgeschlagenen Visualisierungskonzepten, wird derzeit von einigen Forschungsgruppen im Bereich Nanotechnologie eingesetzt.

Abstract

In this cumulative thesis, I describe geometric abstraction as a strategy to create an integrated visualization system for spatial scientific data. The proposed approach creates a multitude of representations of spatial data in two dominant ways. Along the spatiality axis, it gradually removes spatial details and along the visual detail axis, the features are increasingly aggregated and represented by different visual objects. These representations are then integrated into a conceptual abstraction space that enables users to efficiently change the representation to adjust the abstraction level to a task in mind. To enable the expert to perceive correspondence between these representations, controllable animated transitions are provided. Finally, the abstraction space can record user interactions and provides visual indications to guide the expert towards interesting representations for a particular task and data set. Mental models of the experts play a crucial role in the understanding of the abstract representations and are considered in the design of the visualization system to keep the cognitive load low on the user's side. This approach is demonstrated in two distinct fields of placenta research and in silico design of DNA nanostructures. For both fields geometric abstraction facilitates effective visual inspection and modeling. The *Adenita* toolkit, a software for the design of novel DNA nanostructures, implements the proposed visualization concepts. This toolkit, together with the proposed visualization concepts, is currently deployed to several research groups to help them in nanotechnology research.

Contents

Kurzfassung	ix
Abstract	xi
Contents	xiii
I Overview	1
1 Introduction	3
1.1 Problem Description	6
1.2 Research Challenges and Goals	8
1.3 Research Questions	10
1.4 Scope and Contribution	12
1.5 Thesis Structure	16
1.6 Background and Related Work	16
2 Geometric Abstraction for Effective Visualization and Modeling	21
2.1 Using Mental Models for Geometric Abstraction	22
2.2 Distributing Structural Complexity into Multiple Representations	26
2.3 Abstraction Spaces	28
2.4 Integrated Modeling through Geometric Abstractions	39
2.5 Evaluation	43
3 Reflecting Remarks	47
3.1 Impact	47
3.2 Discussion	48
3.3 Conclusion and Future Work	51
3.4 Authorship statement	52
Bibliography	59
Publications	65
	xiii

Part I
Overview

Introduction

I may not have gone where I
intended to go, but I think I have
ended up where I needed to be.

Douglas Adams


In this thesis, I am describing a technological strategy to reform the geometry of spatial structures in their visualization and propose a novel abstraction space concept that guides users towards interesting representations. In the visualization research field, the goal is often to develop a particular set of representations that deals with features relevant for particular tasks. In this thesis, the application-cases require several levels of detail to fulfill user tasks. The underlying idea of this thesis is, hence, to provide expert users with a rich spectrum of visual encodings that can dynamically adapt to a given task in mind. The rising complexity of spatial data requires new ways of simplification. Scientists often study phenomena where many different levels of detail carry information. For example, in architecture, the entities of interest can vary between bricks, walls, rooms and entire buildings. Depending on the task, some features are better observed at some levels than others. Scientific data carry information that is relevant for specific tasks and requires different representations to emphasize these features. Abstraction is applied to reduce the complexity of the data so that only the relevant information remains. The proposed geometric abstraction in this thesis is a pipeline with several steps. In the first step, it deals with the abstraction of scientific data into simplified representations, aiming to reduce the cognitive effort that is needed on the side of the user. Mental models will be considered for the visual encoding of the spatial structures. In the second step, I integrate these representations into a *conceptual abstraction space* that gives the user control to effectively switch and perceive correspondences between visual encodings through animation. In the final step, I augment this abstraction space, which represents the

expert's scope of interaction possibilities, with visual clues to guide the user to potentially relevant views.

There have been great advances and discoveries made in many sciences, where the structural shape is of key importance. In medicine, novel image acquisition and processing techniques lay the groundwork and output spatial data that call for advanced visualization and inspection methods. In the emerging nanotechnology field, structures at the nanoscale have first to be visualized and synthesized with the computer (*in silico design*), before wet lab experiments can be carried out. The need for the synthesis of computational models requires interactive modeling capabilities. Visualization plays a fundamental role in presenting the spatial structures in both fields. With technological advances, we have now the tools at hand to study structural phenomena in a way that was not possible before. But the premise is that data with high structural complexity is presented appropriately, so human perception can process it effectively. The challenge here is to determine the appropriateness and find out at which level of detail the data should be presented to the user. This problem is demonstrated in Figure 1.1 that depicts two common representations of a bio-molecular structure. In the atomistic representation (Figure 1.1a), all the detail can be perceived, but the abstract representation (Figure 1.1b) is far more suitable to understand the high-level features, such as the DNA strands and the protein secondary structure. This example shows why visual abstraction is necessary for structural simplification. In my thesis, I propose a geometric abstraction approach that integrates both, high and low detail levels together and combines their advantages. Especially, the design of nanostructures requires an adaptive approach, since one representation cannot reveal all the rich structural details. The challenge here is to uncover the entire richness of features, while at the same time to transform the data into a more simplified, but comprehensible version of itself. It is not possible to serve both necessities with one representation. But a set of representations that dynamically adapts to the user's needs could solve this. In this case, methods for perceivable change between different abstraction levels have to be introduced to depict the relationship between the representations.

Researchers have developed many specialized methods to explore and analyze structural phenomena. The problem is that the more specialized a tool becomes, the more its specialization deems itself to be utilized for only revealing a particular part of the rich structural characteristics. Multiple specialized visual tools could deal with this challenge. Keeping the user in control of such a transformation of the representation is often an overlooked challenge, especially their relationship to each other. As a result, the missing coherence between tools is unnecessarily taxing the mental processing abilities with tasks that could be otherwise externalized. The idea of this thesis is to transform this cognitive task into a perceptual task, to keep the efforts of interpreting representations low.

The validity of this integrative approach can be demonstrated in case studies of two distinct fields, one being the traditional field of medicine, and the other being the emerging field of nanotechnology. In both cases, the structural characteristics are of tremendous importance, but spatial complexity is also hiding many features of interest. The problems



(a) Concrete: The structure has high visual complexity as all atoms and bonds are depicted. In this representation, high-level features cannot be perceived.

(b) Abstract: The DNA single strands are visualized as tubes. The protein is depicted in its secondary structure using SAMSON's visual model [47].

Figure 1.1: The abstraction of a nucleosome (PDB-ID: 1AOI). Two representations depict the same molecular structure but highlight different features.

in both fields go significantly beyond the basic understanding of shape. Different spatial arrangements are required for an unhindered view of the data. The richness of structural information and the needs of both fields for complexity reduction drives new advances of geometric abstraction in visualization systems. Therefore, this thesis proposes geometric abstraction as a dynamic method for exploring such structural richness.

Based on the two long-term case studies in collaboration with domain scientists with various backgrounds the results for this thesis were obtained. In the first case study, I collaborated with clinical experts and researchers that investigate the in utero placenta. The experts had a particular requirement for the structural simplification of the placenta for assessing fetal health. The initial idea towards geometric abstraction was developed in this case study. The goal was to abstract the spatial properties of the placenta by virtually flattening it. *PlacentaMaps* (**Paper A**) reformats the highly variable shape into simplified 2D representations, enabling expert users a straightforward inspection of this vital organ. Geometric abstraction of spatial properties developed into a more refined concept in a second case study, where the goal was to create a visual tool for the design of DNA nanostructures. As these structures are large and consist of several hundred thousands of atoms, the scene is cluttered with visual details that need carefully designed representations (**Paper B**). The *spatiality* and the *visual details* are considered as the main ways, how structural simplification could be achieved. Together these two ways of abstraction construct a conceptual abstraction space that provides the expert users with a way to navigate to visualizations that suit their tasks optimally (**Paper C**). The developed visual tools are integrated into the *Adenita* [40] toolkit that expert users

adopted to design novel static nanostructures, as well as DNA nanorobots.

1.0.1 Definitions

A *level of detail*, as used in this work describes the degree of complexity in the inspected data and its representation. *Abstraction* follows the definition of Viola and Isenberg [67] and describes the transformation, where only key concepts and details are preserved and visually emphasized. The non-relevant details can be attributed to noise or information that is not important for conveying the intended meaning. The term *geometric abstraction* describes a reduction of shape details in object space, rearranges them in their spatiality and maps the features of interest to the visual details to create effective representations. The abstraction itself is understood as a transformation, while the result of this process is the representation. The two ways of how the structure can be geometrically abstracted (spatiality and visual details) are described as *abstraction axes*. On an abstraction axis, representations are defined at uniform distances, which are also referred here as *abstraction levels*. If the representation is the result of geometric abstraction then the level of detail is referred to as abstraction level. The term *spatiality* refers to spatial dimensions as it refers to the positional features in the data. The term *visual detail* refers to the shape parameters and color of the used geometric primitives for the representation. The two abstraction axes form a 2D coordinate system that defines a representation at each coordinate. This coordinate system is referred to as an abstraction space. The term *spatial data* describes data sets that contain information about physical objects.

1.1 Problem Description

The topic of geometric abstraction is built from the bottom-up based on the research carried out in the two case studies. The case-specific goals are described in detail in the individual publications in Section 1.4. In this section, I will give a domain characterization [46], by summarizing the tasks and describing their requirements for the developed visual encodings.

1.1.1 PlacentaMaps

The first case study was conducted in the *PlacentaMaps* project. The placenta influences fetal health, such as birth weight [43] and neuro-development [48]. Therefore, the idea of the medical experts was to create 2D representations to quickly compare multiple placentas with each other and gain new insights for their research of this vital organ. The tasks of the medical experts are to assess a large number of placentas and to find commonalities and pathologies. They would like to visualize biomarkers, such as overall appearance, shape, attachment area to the uterine wall and other pathological occurrences. The problem is that the placenta is highly versatile in shape, which makes a comparison challenging. Unfortunately, the 3D shape is limiting the ability to analyze and compare multiple in utero placentas, which is crucial in large population studies. A volume rendering would depict the 3D shape without distortion, but a comparison would

be difficult due to the inherent occlusion and necessary interactions. First, the experts want to investigate the fetal side to look for pathological appearances with respect to the umbilical cord insertion and vessels. On the maternal side, they would like to investigate the cotyledons of the placenta. A slice-by-slice view is the most detailed inspection of the placenta tissue, however, this viewing method would be cumbersome when comparing multiple volumes as the slicing would not take the structure of the placenta into account. A necessary requirement was that the approach creates additional views automatically.

Therefore, the goal is to automatically complement the 3D shape with a standardized visual representation of the fetal and maternal side. First, an automated segmentation will extract the placenta from the fetal magnetic resonance images (MRI). A structure-aware slice method extracts a distance field from the segmented mask and creates offset surfaces, which are then flattened into two standardized views of the isolated placenta sides. The main problem of this approach is the distortion of the flattening approach, that could result in 2D images that are hard to relate to the 3D shape. More details on the problem characterization can be found in **Paper A**.

1.1.2 DNA Nanotechnology Visualization

For the majority of my PhD, I worked as a visualization researcher and developer in the *MARA* project [4]. The goal of the project itself was to develop a DNA nanorobot that target pathogens, such as antibiotic-resistant bacteria. As the design process of such a complex structure is mostly computational, my goal was to develop the visual tools enabling the scientists to design these structures. Thereby, visualization is a crucial part as it lays the foundation for subsequent analyses and modeling tasks. In this project, it was more difficult to identify the goals, due to the novelty of the field itself and its challenges. The idea of structural simplification through geometric abstraction has already been initiated in the PlacentaMaps project. This idea was continued and matured in the MARA project. Another ambitious goal was to go beyond a concept and develop a tool that will be released to expert users for designing novel DNA nanostructures. The typical tasks of the domain scientists are described as follows: First, they come up with a conceptual design of the targeted architecture using hand drawing. In this stage, only rough measurements and the high-level geometry, such as the shape and sizes are considered. Usually, some experimentally tested components are used to assemble higher-order entities. In the next step, the experts plan how this target architecture can be created using DNA as a building material. The experts rely here on schematic drawings, where crossovers and the routing of the single strands are well observable. Hereby, the architecture is not only guaranteed by the routing of the single strands, but also the sequences of the nucleotides that determine how the single strands bind to each other to form the characteristic double helix. Once the structure is designed, the sequences of the DNA bases are exported as a text file. One problem here is that the relationship between the applied visual tools is not explicitly shown. For example, the sequences correspond to strands in the DNA schematics, but the representations are not linked in any way. Another problem is the mismatch between the required level of

detail and certain tasks. For example, to design the structure, they have to settle on one level of detail at the beginning. If a high level of detail is chosen, e.g. atoms and bonds, then the expert has the most precise way of modifying a structure. However, having to model the individual atoms would not be feasible if a large structure has to be created that consists of hundreds of thousands of atoms. On the other hand, if a low-level of detail is chosen, then large structures can be easily designed, but precise tasks, such as adding bonds cannot be carried out on this level. The main problem of features that exist on different levels of detail is demonstrated here by the different visual tools that are used to accomplish these tasks. This problem is aggravated by the fact that it is not clear how modeling operations behave across several abstraction levels. Therefore, an integration of the different abstraction levels is crucial for the effective *in silico* design, as described in **Paper B**. The experts also use different layouts for designing structures. For example, in the most popular tool caDNAno [16], the DNA can be only drawn in the 2D layout. Inspection can be performed in a linked view that depicts the 3D architecture. The experts cannot carry out the same operations in 3D, 2D, and 1D, but instead, all operations have to be made in the 2D schematics. For example, in a nanotube, it is important to select those strands that are on the inner side of the tube. This simple task requires the user to select the appropriate elements in 3D. DNA crossovers are connections between the double helices, which can be easily detected in 2D as they are depicted as vertical connections. Finally, the lengths of strands can be easily determined by inspecting the 1D view, where all strands can be linearly aligned. Since there are many different representations and tools in use, it becomes increasingly difficult to understand how the representations relate to each other. The experts are required to mentally match between these representations, which makes the design cumbersome. A more detailed domain characterization is given in **Paper C**.

1.2 Research Challenges and Goals

The main goal is to create multiple abstraction levels that highlight relevant features to serve user tasks and to integrate these representations into a continuous abstraction space. This goal consists of several subgoals that are described in the following:

The abstraction of data by tight collaboration: In a multi-disciplinary setting, a knowledge gap between the visualization researcher and the domain scientists has to be closed before the collaboration can be efficient. The ultimate goal for the visualization system is to solve the target domain tasks. Characterizing the tasks and data in the language of the domain is the first and most crucial step to create appropriate abstractions of the data [46]. First and foremost, the tasks from a high-level to a low-level have to be characterized. At a high-level, the experts could describe a general challenge, such as “I want to be able to view the maternal side and compare it to other placentas.” or a specific low-level task can be formulated, e.g. “I want to identify DNA strands, that are candidates for connection and insert a specific sequence of bases to merge two structures.”. Close collaboration is essential for characterizing the domain and its visualization-related challenges.

Develop an integrated visualization system with various levels of abstraction:

To deal with the variation in tasks, the visualization system has to inevitably produce a set of representations that can be quickly adjusted to the level of detail required. In addition to creating the right visual encodings, the visualization system has to provide order for these representations. The representations can be ordered in a way that enables the user to quickly find the right abstraction level for a given task, for example, going from concrete to abstract. Necessary interaction possibilities have to be provided to the users so that they can quickly adjust the abstraction level to fit a given task. Later, the concept of abstraction axis will be described as the first step to reach this goal. The final goal is to develop the conceptual abstraction space that orders the representations into a controllable space.

Describing a mental model-aware abstraction process: In the visualization field, the concept of visual metaphors is central to make complex information more comprehensible. As Lakoff and Johnson [37] stated, the human conceptual system is fundamentally metaphorical in character. In this context, the tasks go beyond finding visual metaphors, since the challenge is to consider the preexisting mental models of the targeted experts and how they can be interfaced. By characterizing the domain, the mental models could be interfaced. Usually, experts in one field went through similar education and were exposed to the same concepts while learning and building these mental models. Choosing the right visual metaphors is a challenge and is aggravated by the fact that users need to accomplish different tasks and therefore, might require different visualizations. Once the right metaphors for the visual encodings are chosen, the entire visualization concept has still to be created in such a way that the cognitive effort is kept low. In simple terms, it should be self-explanatory and easy to grasp. On the one hand, this concerns the interactions with the visualization system. On the other hand, this concerns how the individual elements in the visualization design are related to each other. Of course, the mental models that the experts apply to interpret the data are difficult to understand for an outsider. However, when learning their workflow and characterizing the domain, the visualization researcher can acquire the domain language and understand the associated mental models better. The experts are familiar with the concepts and especially the tools used in their workflow. In many instances, it makes sense to not re-invent the wheel but instead the familiarity with existing methods should be employed as reference point and incorporated into novel visualization systems in some way. The decision to be embedded in the domain team was crucial.

From cognitive to perceptual task: As the main idea of abstraction is to emphasize particular features of interest at a given moment, multiple representations have to be created that highlight different features. By creating multiple specialized representations, one problem is solved, while another one is introduced. As the isolated representations depict the same underlying data, they are related to each other. For the user, establishing the correspondence between these representations would be a cognitive task. The process of relating these representations in their mind puts an unnecessary cognitive load on the user and thus it should be converted to a perceptual task. Either coordination between

views or transitions show this correspondence explicitly. Effective ways of animation between the representations will be described that enable the user to take control over the pace of the transition.

Guide the user to important representations: The mental models of a user and their effect on the tasks are hidden from the visualization researcher. Trying to understand them by learning the domain of the experts is the first step. While high-level feedback can be provided by the experts, the actual interaction of a user with the visualization system could be an important source to derive additional insights about the user interest in a representation. A more quantifiable way would be to allow expert users to record the history of interactions with the visualization system and the usage of the different abstraction levels. This approach would be helpful to understand the real interest in certain representations for specific tasks. Therefore, I propose a method for recording and visualizing the user behavior and provide feedback to the visualization researcher. At the same time, such visualization could also guide other novice users to important representations of a specific data set for a given task.

1.3 Research Questions

The main intention of visualization is to effectively present information through visual imagery. While a large portion of the human brain is dedicated to processing visual stimuli, one should not disregard the associated cognitive effort when an expert is exposed to a visualization system, especially if it is novel. The cognitive effort might be even larger if there are multiple representations with different visual encodings. While often the complex spatial data is presented through (scientific) visualizations, the expert's task is still to solve a problem by gaining insight into the underlying data. Therefore, an important challenge of developing an applied visualization method for a domain is to keep the added cognitive load low.

In visualization as well as in the domain sciences, highly specialized representations are created to help experts to solve their difficult tasks. In many cases, not one isolated but instead multiple representations are required to gain a comprehensive understanding of the complex data. In medical images, for example, three principal planes are used to provide a view from orthogonal directions into the data. In molecular biology, the data is often hierarchically organized and, at the highest detail level, there are simply too many objects to be visualized individually.

The representations, resulting from geometric abstraction, have to be organized into an abstraction space to be seamlessly linked with each other. The visualization designs are novel to the targeted group, however, it was the best practice to still provide familiar views for the experts to relate between novel and established visualizations. A novel visualization system first requires the user to fully understand it, to unfold its effectiveness. One challenge of a novel visualization system is that the user has to invest time and effort into learning it. Especially, when they have an existing workflow where visual tools are used. One could use a strategy to show the explicit relationships of features between

the old and the novel visualization system. This could promote the understanding of the novel visualizations.

In the following step, the idea is to establish correspondences between all representations in general. While Coordinated Multiple Views (CMVs) [53] is a method that can handle a few representations well, a more elaborate scheme has to deal with numerous views, which is the case in DNA nanotechnology. Controlled transitions can depict the correspondences between representations explicitly. In the end, entire interactive abstraction spaces should be constructed to integrate all visual tools.

Scientists often have to handle data with high structural complexity on the one hand and on the other hand, representations are needed that are highly effective in depicting the desired features of interest. As the sciences advance, the acquired data and the associated tasks often deal with such challenges. As straightforward the problem statement is, all the more difficult and multifaceted it is to develop appropriate solutions. The overall problem can be divided into the following research questions:


- *How can mental models be used in the visual abstraction of spatial scientific data?* The targeted users and their understanding of subjects influence the way how they respond to representations. Geometric abstraction of the data should be applied in such a way, that the mental models are taken into account as demonstrated in Figure 1.2. To create such representations, the understanding of tasks of experts and subsequently their mental models are required. Getting insight into these mental models is not trivial, but essential when abstracting the data.
- *How can a visualization system deal with the variety of features of interest and their representations?* The domain experts are interested in different features of the data, to answer a question. Not all details are important for a task and only selected parts should be conveyed to the user. For example, for one task a particular visual object is of interest, but for another task, the same object could be distracting or even obstructing important features. The challenge of abstraction is to first isolate the varying importance of features on demand to avoid overloading the user with visual complexity. The challenge here is to maximize the expressiveness of single representations and to provide a comprehensive set of views.
- *How to create effective exploration and interaction spaces for abstract representations?* The intention is not to simply create a set of representations, but also to show the user how the visual objects are related. Besides, finding and switching between representations is important to provide a smooth workflow, where the abstraction level always fits the task. Increasing simplification, from concrete towards abstract, is the general direction of geometric abstraction. Hence, the resulting representations can also be ordered by this aspect. The challenge now is to utilize this order of representations and to define a conceptual space that provides the user to quickly explore and find appropriate representations.

- *How can a user quickly control the level of abstraction to adjust to his or her mental model?* Establishing the correspondence between representations might be taxing the user as this is a cognitive process if no visual indication is provided. The visualization system has to provide interactions to convey the correspondence between representations through meaningful transitions. The challenge is to put the user in control of the transition and adjust it to the pace of his or her mental capacity in real-time.
- *Can the visualization and modeling process be optimized by guiding the user to relevant representations?* In the visualization literature, often the final result of a system is presented with some validation by domain experts. It is often not discussed how the feedback can be incorporated back into the redesign of the system. The challenge is that the interactions between user and system have to be quantified. In this thesis, results in the visual quantification of user behavior are shown. A promising aspect is to provide new users with a guide that reveals where the important views are located in a vast abstraction space.
- *How can abstraction help to create visualizations that serve both analysis and modeling tasks?* For DNA nanostructures, the task goes beyond the visualization, but also to synthesize and modify them. The designed visualizations should not only be able to depict the visual features of interest, but also be suited for direct manipulation. The interactive modeling aspect will also be explored in this thesis. Based on this conceptual abstraction space, new ways of applying modeling operations and interactions can be developed.

1.4 Scope and Contribution

This thesis presents the main body of my research work since I started my doctoral studies in October 2015. In two separate case studies in collaboration with medical experts as well as with DNA nanotechnology experts, I developed systems to visualize complex three-dimensional phenomena for the respective users to explore, analyze, and understand the scientific data. Although the projects have been realized in multi-disciplinary settings, the primary contributions of my research were made in the visualization field. The goals were always to find ways to create structural simplifications and thereby, designing specialized, but highly effective visual representations.

From the research and collaboration with domain scientists, several publications resulted, of which three are most coherently connected and, therefore, included in this thesis. They delineate the development of geometric abstraction as the overarching strategy, based on the research in the case studies and the exposure to the challenges in the respective domains. The main contributions of this thesis are from the *PlacentaMaps* [44] and the *MARA* [4] project and their resulting publications. Figure 1.3 shows how these publications fit together in the overall topic of geometric abstraction. The following publications make up this cumulative thesis:



(a) Virtual flattening of the in utero placenta. (b) The abstraction of a 3D DNA nanobottle.

Figure 1.2: The effect of mental models on the visualization design. In both cases, a 2D representation is generated, but using different previous knowledge of the domain scientists. In (a) the mental model is based on previous knowledge about the real-world behavior of the placenta after birth. (b) shows the complex 3D shape of a DNA nanostructure that is abstracted into a 2D representation based on familiar DNA schematics.

Paper A - Placenta Maps: In Utero Placental Health Assessment of the Human Fetus. Haichao Miao, Gabriel Mistelbauer, Alexey Karimov, Amir Alansary, Alice Davidson, David F.A. Lloyd, Mellisa Damodaram, Lisa Story, Jana Hutter, Joseph V. Hajnal, Mary Rutherford, Bernhard Preim, Bernhard Kainz, M. Eduard Gröller. *IEEE Transactions on Visualization and Computer Graphics*. 23(6):1612-1623, 2017. [Best Paper Honorable Mention at the IEEE Pacific Vis Conference 2017]

Paper B - Multiscale Visualization and Scale-adaptive Modification of DNA Nanostructures. Haichao Miao, Elisa De Llano, Johannes Sorger, Yasaman Ahmadi, Tadija Kekic, Tobias Isenberg, M. Eduard Gröller, Ivan Barišić, and Ivan Viola. *IEEE Transactions on Visualization and Computer Graphics*. 24(1):1014-1024, 2018.

Paper C - DimSUM: Dimension and Scale Unifying Map for Visual Abstraction of DNA Origami Structures. Haichao Miao, Elisa De Llano, Tobias Isenberg, M. Eduard Gröller, Ivan Barišić, and Ivan Viola. *Computer Graphics Forum*. 37(3): 403-413, 2018.

The overall contributions of this cumulative thesis are:

- The concept of geometric abstraction for two domains using multiple visualizations

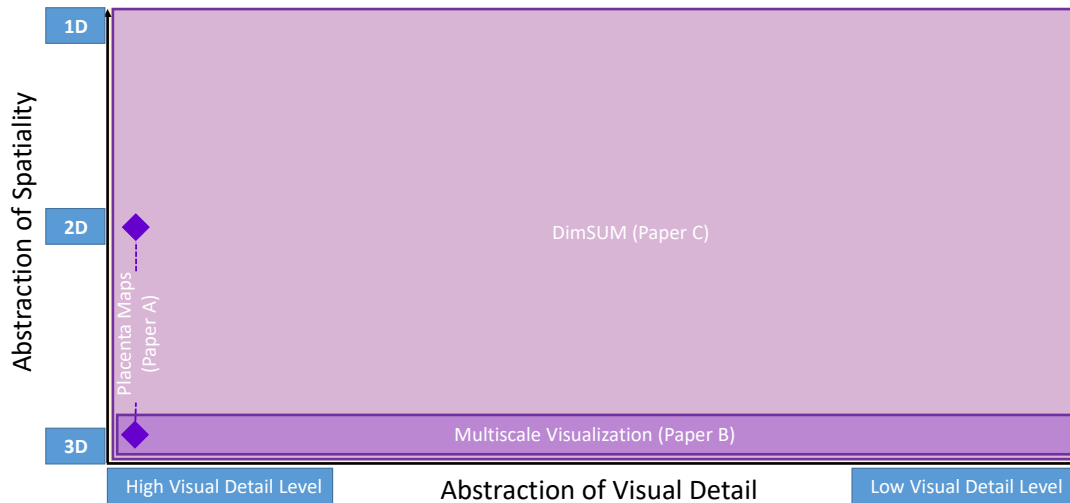


Figure 1.3: The proposed geometric abstraction method creates a conceptual abstraction space. This space is spanned by the spatiality and the visual detail axes. With every publication, the coverage of this conceptual space could be increased. The dashed line in the PlacentaMaps approach (**Paper A**) indicates that there is no transition defined, hence there is no continuum. The multiscale visualization concept from **Paper B** provides transitions across 3D representations and the DimSUM approach in **Paper C** covers the entire abstraction space.

to highlight various features in the data.

- A novel geometric transition concept to connect multiple visualizations into a conceptual abstraction space.
- A description of a visualization design process aiming to limit the cognitive load when working with complex spatial data.
- An abstraction process that facilitates interactive modeling across different abstraction levels.
- A reflection on geometric abstraction and the application to other domains.

The contributions of **Paper A** are:

- An automatic approach to split the in utero placenta into the fetal and maternal side.

- A well-defined structure-aware slicing method that enables experts to inspect the interior and exterior of the placenta.
- A visual mapping of the fetal and maternal side to standardized visual representations, allowing clinical experts a concise, comparative visual assessment.

The contributions of **Paper B** are:

- The domain characterization for DNA nanotechnology visualization.
- A novel continuous, multiscale visualization concept that reduces the complexity of the DNA nanostructures through multiple consecutive detailed and abstracted representations. The complexity of these structures is distributed among several visual detail levels.
- A 1D abstraction space that integrates the consecutive representations using seamless, cohesive transition between them.
- A set of operations that facilitate consistent manipulation of DNA nanostructures that is scale-adaptive.
- A general visual representation of DNA nanostructures that is able to depict the result according to different DNA design concepts.

The contributions of **Paper C** are:

- The discussion of abstraction in illustrative visualization is further advanced. A 2D interactive abstraction space that integrates novel and existing DNA nanostructure visualizations.
- Using geometric transitions, the correspondence between established and novel representations can be perceived by domain scientists.
- The proposed concept highlights information along two abstraction axes, the visual details, and spatiality, including 3D, 2D, and 1D representations. The spatiality is an important abstraction axis that goes from a concrete shape depiction to a successfully simpler representation to gradually abstract the alignment of visual objects.
- A number of specific user interactions are implemented to support the navigation of the abstraction space as well as the analysis of task-specific interaction behavior, including abstraction axis snapping, guiding (heat-)maps, and saving viewpoints.

1.5 Thesis Structure

The thesis is divided into two parts. In the overview chapter, the overall strategy of geometric abstraction for visualization and modeling is delineated. This topic provides a contextual link and introduction to the publications **Paper A**, **Paper B**, and **Paper C**. In this chapter, also the impact of the proposed method and its application in the use cases are discussed. In the second chapter, the individual peer-reviewed publications are included that make up this thesis.

1.6 Background and Related Work

The domain-related terminology, the background knowledge, as well as the related work are outlined in this section. This thesis is related to abstraction, mental models, animated transitions, and general visualizations for the respective domains. A more detailed description of the related work can be found in the Publications chapter.

Domain Background: The DNA nanostructure as well as the in utero placenta are both 3D structures that need simplification through abstraction of different aspects in the data. The commonality of a placenta and a nanostructure is their need for structural simplification into 2D representations. While the placenta is abstracted through a virtual flattening method, the nanostructure is mapped to a 2D schematic layout. In both cases, the challenge goes beyond creating these 2D representations. Another common challenge is to relate the 2D with the 3D representations to depict how they correspond. This form of complexity reduction of the placenta is a visualization approach that is often employed to flatten anatomical structures [33]. The placenta is a highly versatile organ that has two sides, one facing to the uterine wall (maternal side) and the other one facing the fetus (fetal side)[6]. On both sides, some details might be of clinical interest. It has been long known that the pathological findings in the placenta cause preterm labor [3]. There is no standardized, 2D representation of the maternal and fetal side that would allow clinical researchers an efficient analysis and comparison of the placental structure. One consequence is that large population studies are difficult. A DNA nanostructure distinguishes itself from the placenta in the acquisition. While the fetal magnetic resonance imaging data is acquired from pregnant women, a DNA nanostructure is a model that is designed in silico. As the name suggests, these nanostructures use DNA as the main building material. The basic idea is to program the sequences of DNA single strands so that they will self-assemble into a pre-determined architecture. Using the DNA origami method [57], experts exploit base-pair complementarity [68] to create pre-determined shapes when single strands form the characteristic double helical structure. This self-assembly is achieved by programming the sequences of the single strands. The principle is to use a long strand called *scaffold* that is folded together by short, complementary *staple strands*. The double helices are locked into position by single strands that cross over between double strands. The caDNAno tool [16] supports the fast prototyping of scaffolded DNA origami structures using 2D schematics. This tool provides the user with CMVs to enable his or her to design the target geometry in 2D

and the inspection in 3D. There are automatic and semi-automatic methods, such as DAEDALUS [65] and vHelix [7] that take a polyhedral structure as an input and output the sequences for the experiments. As the field is young and rapidly growing, there is a need for visualizing these structures effectively, so that interactive modeling can be facilitated.

Visual abstraction: Viola and Isenberg [67] formalized abstraction in the context of illustrative visualization. Visual abstraction differs from the conventional visualization design as it transforms the data into a representation by removing details. The authors refer to such an organization of representations as an axis of abstraction. This terminology is adopted in the scope of this thesis as well. Low- and high-level abstractions were also discussed by Rautek et al. [51] in the context of illustrative visualization. While low-level visual abstraction concerns the visual styles, high-level abstraction techniques exploit knowledge about the features and suppress and emphasize them according to their importance. The geometric abstraction in this thesis corresponds to a high-level abstraction.

Abstraction along the aspect of visual details: To some extent, the idea of geometric abstraction along the visual details is rooted in the illustrative rendering field. For expressive visualization and presentation of scientific illustrations, often non-photorealistic rendering methods are employed. The object is thereby not only abstracted using different (virtual) depiction techniques, such as hatching [50, 71], stippling [41, 42], and toon-shading [13], but features are also suppressed and occluding objects are moved aside to depict internal structures for illustrative purposes [17, 66]. A cartoon representation is often employed for molecular visualization. In biological visualization, different techniques are used to represent molecular features. The abstraction of molecular surfaces has been investigated in the past decades [32, 12]. Traditionally, researchers distinguish between several techniques that can highlight different aspects of the molecules. For atomistic models, the atoms and bonds can be represented with simple spheres or cylinders. In structural biology, the different surface models are used to represent the surface properties of a molecule.

Gradually reducing the complexity of the models is done if large molecular systems have to be simulated. In molecular dynamics simulations, the systems consist of a large number of atoms, which are computationally too expensive to simulate. The problem of these large-scale systems is dealt with coarse-grained models that represent the complex system in a reduced manner [30]. This coarse-graining is analogous to the construction of visual detail levels as demonstrated in **Paper B**, where the representation are increasingly abstracted. In DNA nanotechnology, the important features exist on several visual detail levels.

Abstraction along the aspect of spatiality: There are many visualization methods that abstracts the spatiality aspects in data. An important family of methods that abstract 3D into 2D shapes are projections. Hence they are regarded as an abstraction of spatiality. These methods are commonly used in engineering and architecture. A traditional field is map projections that map the earth onto a plane. Depending on

the preserved features (for example angle and area) different map projections have been developed over the past millennia. In medical visualization, often some key shape characteristics can be used to flatten organs and vessels into a 2D plane [33]. Within a 2D plane, the layout of the elements can be changed [69]. Metro maps are schematic drawings, that trade the exact geographical location for highlighting other features. In the biological domain, the primary, secondary, and tertiary structures are abstract depictions at different spatiality levels.

Correspondence between representations: A natural candidate of coordinating representations are CMVs. They combine and link different views to increase the insight into the data by showing local correspondences. Using brushing and linking, relationships and connections can be discovered that would be hidden if no coordination between views happens [10]. A more critical view on CMVs was given by Andrienko and Andrienko [2]. The amount of data poses a scalability problem for CMVs. If the features at the highest detail level in the data are represented, then the screen space would be a limiting factor. In **Paper B** and **Paper C**, it is the amount of data (number of atoms) that inhibit this problem. The authors also state that there is a sheer number of existing visualization methods and that it is difficult to find a combination of suitable and sufficient views. This issue is addressed in this work by providing a large set of abstraction levels that can be explored to find a suitable representation. Another limitation of CMVs are the small number of views that can fit on the screen. Due to the limitations and new challenges in the two case studies, controlled transitions are employed. To deal with varying tasks, one established way is to employ complementary representations, such as CMVs, to provide different ways of viewing. The challenge is to convert a cognitively demanding task into a perceptual task [54]. Portions of cognitive workings should be shifted into perceptual processing. The associated cognitive efforts can be reduced if the representations are matching the mental ones. Existing mental models have to be understood first to contribute to the design of a novel visualization system. Animation can be applied to help the user to understand the correspondence between representations. Compelling evidence for the efficacy of animated transitions in statistical data visualizations has been shown by Heer and Robertson [26]. The efficacy of animations has also been discussed by Tversky et al. [63]. From the cognitive science point of view, animation is regarded to be less effective to convey complex systems than static graphics. The most relevant insight of the authors is that interactivity can overcome the drawbacks of animations in general. Hence, the idea in this thesis is to give the user control over the animation. Fisher [18] stated that animation can be useful to show the logic behind an idea by depicting the intermediate steps. The author also concluded that animation imposes a burden of complexity and it should pay off. In **Paper A** the animated transition between views was not necessary, as the correspondence between views was easy to establish due to the small number of views. In **Paper C** the system transitions between 3D, 2D, and 1D. The 3D representation is easily understood, as it is a model approximating the real shape. The 2D and 1D representations are specialized layouts that require expert knowledge of DNA structures. In this case, the transition was necessary to show the intermediate steps between the 3D and the other views.

The visual detail and spatiality axes can be united into the overarching concept, which is referred to as an abstraction space. Van der Zwan et al. [64] illustrates the combination of abstraction dimensions into an abstraction space. Multiple abstraction axes describe the various aspects in the data with increasing levels of abstraction. The concept of an abstraction space that enables the user to control the representation has also been proposed by Mohammed et al. [45].

Visualizing domain-specific spatial data: A thorough discussion on the related work concerning the visualization of the structural data for a specific domain is given in **Paper A**, **Paper B**, and **Paper C**. The spatial qualities in 3D data are often of primary interest, which makes the representation on typical 2D display media a common challenge. For medical data, projections [58] and reformations [28, 34] are used to highlight various features by arranging the visual elements according to an established standard. When dealing with both, biological and medical data, the highly specialized representations only convey partial aspects, but intentionally disregard details that are not relevant for a particular task.

Interactive modeling: A primary purpose of the visualization done in this work is to facilitate effective interactive modeling. Interactivity is thus a relevant aspect in this thesis. Yi et al. [70] reviews interaction in information visualization. They propose seven different categories of interaction techniques based on the user intent rather than low-level techniques. The proposed interaction technique to control the abstraction level would fit into the *Abstract/Elaborate* category.

The proposed abstraction space as an interactive element is implemented as a panel that can be moved around on the screen. Marking menus [36] is an effective technique to activate operations based on a particular position of the mouse. The abstraction space could be extended as a context-dependent menu. In the computer-aided design field, methods [56] have been proposed to design objects using several models. For the modeling across several abstraction levels, previous work suggested using a reference grid as view context for sets of complex geometry [21]. Glueck et al. [22] presents a continuous multiscale navigation approach for client-server visualization system. The system creates multiple levels of detail versions of the data to provide progressive loading.

Mental models: A theoretic foundation for the understanding of mental models has been laid by the field of cognitive science. Sweller [61] studied cognitive load in domain-specific knowledge-learning. Cognitive load refers to the effort put on short-term memory, which can be seen as limited mental capacity. It is self-evident that the interpretation of a visual encoding should tax the expert's mental capacity. Therefore, visualization systems should utilize the existing mental models of the experts as much as possible. Previous knowledge and familiar concepts can be used for keeping the cognitive footprint low. The preconceived knowledge is structured into concepts and their relationships that represent real-world systems. These concepts are referred to as mental models [19]. The author also states that nobody images all the world in their heads, but only selected concepts. This fits well to the idea of using abstraction to create specialized representations of selected sets of features in the data. Previous work in information visualization [52] already

argues that the comprehension of information graphics is promoted by existing schematic structures. As we use different mental models to understand the selected concepts, I assume that the explicit depiction of the correspondences between these models can further decrease the cognitive footprint.

Creating specialized and effective visualizations requires us to understand the domain scientists and ultimately understand their previous experience with existing ways of representation. The relevance of the interplay between mental models and visualizations has been examined and found to be a key aspect in cognitive offloading [38]. The formal work about this interplay and the application-driven nature of the case studies further support the underlying importance of mental models for abstracting scientific data. For visualization purposes, the preserved features have to be visually encoded to create a specific representation. This (visual) model is the external representation to the experts and has to be perceived and understood. To make the external model effective, an important aspect is to take the mental models of the users into account. Liu and Stasko [38] have already examined the interplay between the internal (mental) models and external representations for information visualization. They define a mental model as functional analog representation to an external interactive visualization system that preserves structural and behavioral properties. Their discussion of the effect of mental models on the visualization design and evaluation is relevant for scientific visualization as well. Trafton et al. [62] described how scientists connect the internal and external representation in scientific visualization. They found that scientists perform spatial transformations on external scientific visualizations and align them with the internal mental representations. The authors describe the spatial transformation as a cognitive process on either an external or internal representation (e.g. rotation). The authors believe that this transformation is represented by a visual component and a spatial component.

Geometric Abstraction for Effective Visualization and Modeling

Animation can explain whatever
the mind of man can conceive.

Walt Disney

In scientific visualization, many representations have been developed that depict certain features in spatial data that serves the exploration and analysis tasks of experts. Spatial data describe 3D, 2D and 1D objects. Their structural properties are often the subject of analysis. To depict the features of interest and communicate the intended meaning to experts, the spatial data have to be reduced in complexity. The occlusion problem is an inherent challenge in 3D visualization that has been addressed using different strategies [17].

A key strategy to reduce the spatial complexity of the data is to visually encode only the features of interest in the data and in this way create specialized representations. I will describe ways and sources that can inspire such an abstraction process. In the next step, a linear order of the representations has to be established. The proposed continuous abstraction enables users to control the pace of the transitioned animation between two representations. The control over the speed of the animation is important, as it allows the user to perceive and understand the transition at any pace. Finally, a novel way to visualize the user's interactions with the system is introduced. The visual quantification of user behavior can reinforce or reject initial design decisions and hypotheses. For the user, it acts as guidance towards relevant views, based on the behavior of other users.

The complex spatial data contains details that are not always relevant to a specific purpose. Dealing with the structural data, the spatial character is of interest, but not all of it. The inspiration for such an idea comes from established structural representations of DNA in the biology domain. A DNA molecule can be described as a sequence of bases (primary structure), as the interaction between these bases (secondary structure), or as the arrangement of the individual atoms in three-dimensional space (tertiary structure). While all three structures represent the same DNA molecule, they depict different features of the same real-world phenomenon. As real-world objects are often complex, we have to sacrifice parts in the data to make selected concepts understandable. From these observations, the idea of geometric abstraction developed. It describes the idea of a transformation process that converts spatial data into multiple simplified representations and in this way highlights features, depending on the task in mind.

2.1 Using Mental Models for Geometric Abstraction

A user's mental capacity should be primarily used to solve the tasks and not to decode the visualization. From this proposition, there are three ways, how mental models can impact the effectiveness of the visualizations. First, the visualization researcher has to understand the familiar concepts of the targeted user as much as possible and then incorporate them in the new visual encoding. The users are observed and interviewed to characterize the domain and to derive the requirements on the visualization tool [46]. Both, the internal mental model and external representation are simplifications of real-world systems that explain general or certain subjects.

Scientists often use their internal mental representations to help them solve problems and reason when using scientific visualizations [62]. The scientists have to connect the external with the internal representation. A close match between the expected internal representation and the exposed external visualization should be aimed for. A mismatch could cause spatial transformations, which in the worst case could potentially lead to misinterpretation. Another pitfall could be that the capability of the targeted group to interpret novel visualizations is overestimated. Considering their mental models could be an important step in designing representations that can be interpreted better. Furthermore, interpretation of the novel representation could be eased by animating the transition to the old, familiar representation. Second, to accomplish a high adoption rate of the visual tools, the abstraction levels should align with the tasks. One should be aware that the novel visualization could be regarded as more tedious to fulfill the actual tasks. Especially if there exist previous visual tools that experts have already learned to use. An additional tool would require them to invest time to learn it. Research often tends to be driven by novelty. One potential pitfall in the visualization research could be that the established concepts from existing workflows and their used visual elements are overlooked and deemed as inferior per default. Presenting the user with an explicit depiction of the relationships between the existing and novel visualizations could increase the comprehension of the latter. The existing views reveal much about the existing mental models, and thus they should be considered in the abstraction process.

Third, establishing the correspondence between representations should be a perceptual task. In the course of this work, the explicit transitions through animation have shown to be very effective in depicting the correspondence between representations and their features. While it is not possible to look into the minds of the expert and to fully understand the mental models, it is feasible to infer them from different sources as described in the following.

Domain literature: To gain a better understanding of the mental models, it might be interesting to investigate the literature they used to study. Especially, scientific illustrations in textbooks play an important role in forming these mental models, as scientists learn and study them during their education. An illustrator intends to make expressive drawings to explain a certain phenomenon. In scientific illustrations, the object of interest is often manipulated to show those features that should be emphasized. Often, the illustrator deforms and distorts the geometric properties to reveal and depict the intended message. The resulting illustrations often deviate from the real shape and behavior of the objects as a trade-off for a more effective explanation of the intended meaning. The scientific illustrator studied the topic and is well aware of the information that should be communicated and hence emphasizes it. Another important source are figures in publications. As the publication process is highly competitive, the figures are powerful explanations, where a lot of knowledge is condensed and the concepts are explained clearly. The illustrations from the publication in the respective fields could also reveal how the experts imagine a representation and in the end, what their mental model might look like. However, one should be aware that these illustrations are only static. The design of the multiscale visualization concept in **Paper B** was inspired by many depictions found in publications of the domain, e.g. the double-strand representation as a cylinder is inspired by Castro et al. [11], the single strand depiction as differently colored tubes with a direction is mainly inspired by the original work on DNA origami [57] and the work of Benson et al. [7]. These depictions are commonly used in domain literature. The usage of the DNA schematics for the 2D representations (**Paper C**) is inspired by the figures used in many publications of the domain [16, 15, 57, 11].

I studied domain literature to derive the scales in the multiscale visualization concept (**Paper B**), as shown in Figure 2.1. The scales are designed to be very familiar to established and well-known representations of DNA nanostructures.

Existing workflows and views: Before starting a collaboration with the domain scientists, they might already have visual tools in use. The importance of the existing tools should not be overlooked, as they can determine how well the new visualization can be understood and how they go beyond their current approach. From the collaboration with domain scientists, it has been revealed that incorporating familiar corresponding views is very beneficial for introducing novel representations. This approach makes the advantages and features of the new system more apparent and at the same time, reveals the short-comings of the existing views. Incorporating familiar concepts into the novel visualization is obviously not a straightforward decision. One should first investigate and question the existing methods and then implement the selected concepts and ideas that

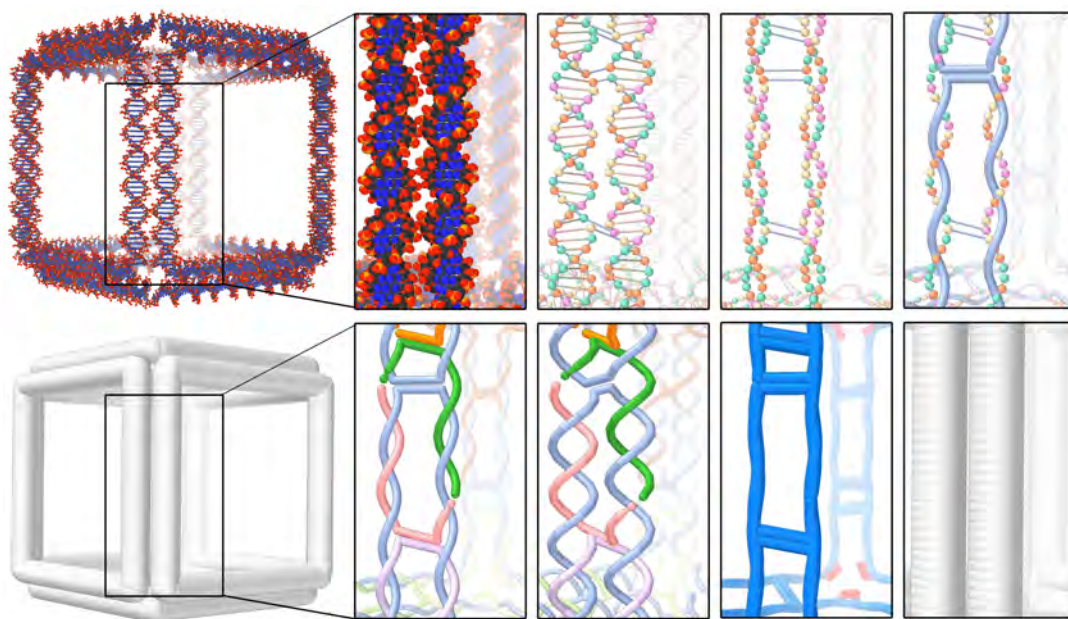


Figure 2.1: The figure shows a nanoscale cube generated with the approach by Veneziano et al. [65]. It is visualized at different visual detail levels. The visual details vary and the features, such as atoms, nucleotides, single strands, and double strands are encoded differently.

are meaningful.

The existing workflows are not restricted to software tools. For example, in **Paper A**, it was not clear for quite some time how to create a flattened two-dimensional representation as there were no existing software tools. In this case study, the workflow of the after-birth investigation of the placenta was a source of inspiration. After birth, the placenta is usually placed on a table where it assumes a round flat shape, with one side facing up and the other down. Figure 2.2a shows the after birth placenta that is very familiar to the medical experts. The idea of splitting the placenta into the maternal and fetal side before flattening reduces the distortion and in this way makes the flattened structure more comprehensible.

Paper C proposes corresponding visualizations in 3D, 2D, and 1D for DNA nanostructures as shown in Figure 2.3. To create a new system that is familiar to the domain experts, I spent much of the time working with the tools used in the current workflow and reproduced their tasks to understand the current drawbacks. One key observation is that the experts need to model in all spatial dimensions, which was not possible with the existing tools, such as caDNAno [16]. Unfortunately, there are no real-world examples of 2D layouts of DNA nanostructures that could be used for inspiration, as in the PlacentaMaps project. Instead of developing entirely new representations, I based the novel visualization concept on familiar 2D DNA schematics [57, 16]. This strategy




Figure 2.2: The comprehension of the virtual flattening approach is promoted by basing it on the deformation behavior of the actual after-birth placenta, which is highly familiar to the experts. (a) and (b) show photos of the deformation into a flat circular shape. (c) depicts the virtual splitting approach into the fetal and maternal side. (d) and (e) depict the resulting representations of the virtual flattening approach.

eased the way for the novel 3D and 1D views, especially in combination with controlled transitions and their animations.

Hand-drawing and sketches: Another effective way to understand the experts' needs is to let them explain the problems and tasks through hand-drawn examples. In the course of my PhD, this method has shown to be effective in discussing the needs of the scientists. The biggest advantage of this method is that hand-drawing is not constraint by the limits of existing tools and therefore, the expert can be more imaginative. As a consequence, novel ideas could be created that are not just an increment of existing tools. One should be aware that this approach might be limited by a user's drawing capabilities. For instance, three-dimensional structures are more difficult to draw and

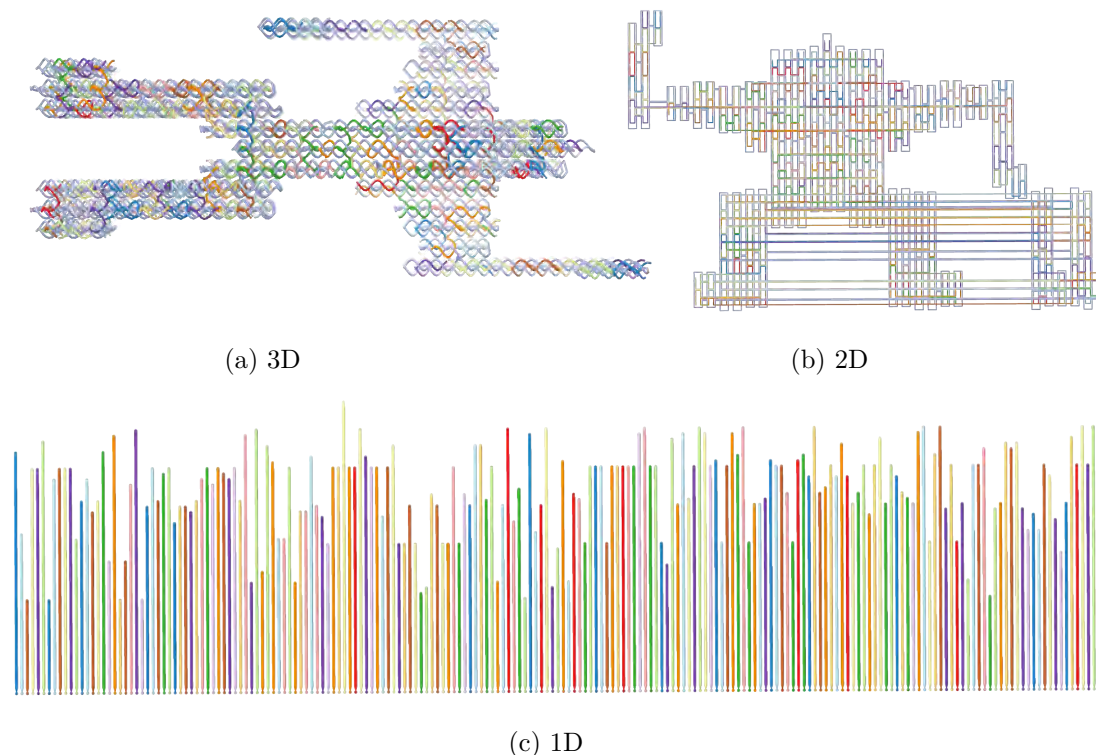


Figure 2.3: A nanorobot [11] is represented in different spatiality levels, from a 3D shape to a 2D schematic layout to 1D sequences.

the expert might rather create two-dimensional depictions. But still, the unconstrained way of explaining outweighs this limitation. Using hand-drawing has been helpful in all three papers (**Paper A**, **Paper B**, **Paper C**) at early design stages, where ideas were sketched out and later when improvements were suggested. Another interesting observation is that the nanostructure designers and users of visualization systems create hand-drawn drafts before they model the structure in the Adenita toolkit. The experts noted that drawing on paper is fast and multiple concepts can be sketched quickly. The interplay between sketching and the visualization system opens up exciting research questions for the future.

2.2 Distributing Structural Complexity into Multiple Representations

One aspect of using geometric abstraction is to design a representation that maximizes the effectiveness in showing only certain features. As a consequence, multiple representations are created along with the abstraction process. This idea distributes the structural

complexity into multiple representations. Only those features that are relevant for a specific task will be presented to the user at a given moment, while non-relevant details are suppressed. The representations are ultimately depicting different, if not mutually exclusive, features in the data. The fewer distracting features, the better is the comprehensibility. Consequently, the cognitive load can be reduced.

The motivation for this approach is that the user cannot process and observe all the details of the data at once. This approach is heavily applied in the case study of DNA nanotechnology. Although DNA nanostructures consist of atoms, the expert thinks in more abstracted representations depending on the task, as demonstrated in **Paper B**. The same is also true in medical visualization, where usually different properties of the organs are investigated. Sometimes, even the shape has to be depicted in an abstracted form to reveal the features of interest. In **Paper A**, the 3D in utero shape of the placenta is complemented with a standardized 2D flattened structure, as shown in Figure 2.2. This idea of complexity reduction through multiple complementary representations is more general than the described use cases. For example, in architecture, different abstraction levels are deeply embedded in the workflow. Construction plans depict the 2D layout of structural elements, but not the individual bricks. A construction worker can derive the brick placement from these plans, however, an architect plans the building in terms of higher-level entities, such as walls. For a city planner, the detail level changes to entire buildings and their relationships. The complexity reduction happens here not only with visual details, but also spatiality, as the plans are usually drawn in 2D. While 2D plans are appreciated for their exact measurements, the 3D visualizations are often used to convey a realistic impression.

The consequence of this approach is that along the abstraction many representations are created to serve particular tasks. An argument for this approach can be made by the way how humans perceive visualizations: We are only able to process a limited amount of information, without overloading our cognitive abilities. Our brain has a limited capacity for processing information and can filter out unwanted information [14]. A key function of geometric abstraction is thus to prepare the data to reduce the unwanted information in the visualization so that the selection process is reduced.

In **Paper B**, I first examined the abstraction levels that the domain scientists used to describe tasks. For example, to explain the task of connecting DNA strands, domain scientists draw tube-like structures or lines to represent the single strands and not the atoms. Although the actual connection happens through atomic bonds, it would not be feasible to carry out such an operation on a high visual detail level due to the visual clutter. On the other hand, the atomic resolution provides a high degree of visual detail that can be used for fine-tuned operations. Due to the task diversity, the structural complexity of the DNA nanostructure is distributed across several abstraction levels that are familiar to the domain experts. Examples are shown in Figure 2.1. Each representation depicts certain structural features of the DNA without other details. The effectiveness results from the simplified representation of the features of interest without distracting details. Another advantage is that the domain scientists can operate on the

abstraction level that fits their tasks the best. This way of abstraction reduces the *visual details* and effects to the number of visual elements and the geometric primitives that represent the underlying data. For example, biological structures have a hierarchy and consist of many atoms at the highest detail level. Representing each atom with a visual element, a sphere, would result in a high visual detail level. Values, such as the position of grouped atoms can be aggregated to reduce the number of elements needed for the representation. If only the aggregated value is represented by a visual element, the visual detail level will decrease. The challenge is to decide at which visual detail level to present the data.

Visual complexity cannot only be reduced in by decreasing visual details but also *spatiality*. This way of abstraction refers to changes in spatial positions of the visual elements when the re-positioning reduces spatial dimensions. With scientific data, the structure is often of primary interest. A 2D representation can depict the details without the problems of 3D, but comes with the cost of losing spatial clues.

In **Paper A**, the abstraction of spatiality was not only done between the 3D and the 2D representation, but also by splitting the placenta first and then flattening the two sides individually. An advantage of the 2D images of the fetal and maternal side is the easy comparability with each other. Depending on the visualization method, a considerable number of interactions are required to inspect a volume. Comparing multiple volumes would further increase the number of interactions, making the comparison a time-consuming process. Besides, the variability, such as the location and shape of the placenta, makes a side-by-side comparison of the volume cognitively straining. The virtual splitting and flattening of the placenta create 2D images that make the comparison relatively straightforward. Regularities, as well as anomalies, could be discovered by looking at many placentas, which is made possible by the abstraction of spatiality. Still, 3D volume rendering has also the advantage of preserving the spatial positions. Hence, I decided to use CMVs to provide both 3D and 2D views to the user. This visualization concept supports the idea of distributing complexity into multiple complementary representations. In **Paper C**, the spatiality is also reduced through complementary 3D, 2D, and 1D representations as shown in Figure 2.3. Each representation provides a different view of the structural details. The 3D view shows the spatial model, the 2D view depicts the DNA crossovers and the bound single strands better and the 1D view linearly aligns all DNA strands.

2.3 Abstraction Spaces

As one representation is designed to depict only certain features in a highly effective way, it can likely happen that, for a single task, several representations are required to provide the user with all the necessary features. It is common that biologists use molecular visualization tools such as VMD [27] or Maestro [59] and switch between several representations to see different features in molecules. Usually, no transition is depicted when switching to another representation, therefore, it is not shown how the features

are related. As the abstraction approach in this thesis gradually removes details, the representations can be ordered by the number of details removed. The result of geometric abstraction are representations that can be organized into a space. This abstraction space is a conceptual construct spanned by the abstraction axes. In this abstraction space a position defines the representation. On the one hand, this method provides an overview of all representations created by geometric abstractions. On the other hand, it acts as an interaction element that enables the user an effective switch and transition between representations to quickly find suitable visualizations.

2.3.1 Geometric Abstraction

In geometric abstraction, the details can be removed in two ways. First, the features can be mapped to visual objects with different detail levels. This approach reduces the visual detail and represents higher-level features. For example, the strand connectivity of DNA is a higher-level feature. To show this feature, the atoms have to be abstracted away as depicted in Figure 2.4. Second, the spatiality can be reduced by removing distracting spatial features as shown in Figure 2.5.




Figure 2.4: An example of geometric abstraction along the visual detail axis of a DNA double helix (PDB-ID: 1BNA). The visual detail level decreases and the shape becomes more and more coarse until the entire double helix is abstracted into an elongated capsule.

2.3.2 Abstraction Axes

While it seems quite straightforward to remove undesired details in scientific data, it is a challenge to represent the preserved features. Viola and Isenberg’s [67] investigation




Figure 2.5: Geometric abstraction of a DNA double helix (PDB-ID: 1BNA) along the spatiality axis. Spatial characteristics can be increasingly removed to emphasize certain features, such as the chemical skeleton or primary structure.

Along the Spatiality Axis	Along the Visual Detail Axis
Map Projection	Glyph-based Visualization [9]
Anatomical Flattening [33]	Molecular Structure Abstraction [32, 12]
Metro Map Layouting [69]	Illustrative Visualization [50, 41, 66]

Table 2.1: Different approaches that geometrically abstract spatial data.

defines several ways in which the data can be reduced as abstraction axes.

Visual detail axis: The complexity in scientific data often arises from a large number of visual objects, their positions, and relationships to each other. Especially, biological data with an inherent hierarchy have such complexity, as there exist many individual elements if the structure is large. Visually encoding all elements individually would not be feasible to depict high-level features, as shown in Figure 1.1. The visual detail level should depict those features the user is interested in, as shown in Figure 2.4. Aggregation can be typically performed to reduce the visual elements and therefore to lower the visual detail level, as demonstrated in **Paper B**. An example where the visual detail level is gradually decreased is shown in Figure 2.1. The highest visual detail level of atoms are depicted by the balls-and-sticks representation. In the first abstraction step, the atoms of a nucleotide are aggregated to an average position. In this way, the atoms are abstracted away and the entire nucleotide is represented by a single bead, instead of several balls and sticks, hence the abstraction is reducing the visual details. Table 2.1 lists representative methods that abstract the visual details. Glyph-based visualizations [9] visually encode the features using a collection of visual objects. With increasing abstraction level, the aggregation of data enables higher level features to be depicted. Therefore, also the visual details will be reduced. In illustrative visualization [50, 41, 23], the intention

is to explain certain phenomena. The focused features are depicted with high visual detail, while the contextual structures are depicted with low visual detail. In molecular visualization [31], bond-centric models depict the full atomistic details, and surface-based methods show different surface properties, hence lower visual details.

Spatiality axis: As scientific visualization often depicts physical entities, the spatial structure is often the subject of investigation. The abstraction along the visual detail axis already creates visual encodings at different hierarchy levels, but does not rearrange the visual elements at different positions in accordance with a certain layout. As typical display media are two-dimensional, such as monitors and paper, there already exist many representations that lay out the visual elements in 2D space. Abstraction along the spatiality axis comes with losing the actual positions of elements and bears the risk of creating representations that have no spatial reference. Therefore, the advantages of such an abstraction need to outweigh the disadvantages.

In **Paper A** the PlacentaMaps approach abstracts a segmented placenta in a 3D volume into 2D images, showing the flattened maternal and fetal sides. The third dimension is abstracted away by mimicking the real-world flattening behavior of the placenta after birth. As shown in Figure 2.6, the 3D view provides a non-distorted view on the shape, location, and structure. The 2D representations provide a complementary view on the two sides of the placenta while sacrificing unimportant spatial features, such as the location and the actual 3D shape. The flattened structures provide advantages, such as simple comparability. In **Paper C**, the abstraction gradually removes spatiality from a 3D nanostructure. Using animation the transition between 3D shape to 2D schematics, and 1D linear alignment of bases can be depicted. In 1D, the spatial reference is further reduced and only neighborhood information is preserved (sequence of nucleotides).

In Table 2.1, representative visualization methods are categorized according to the proposed abstraction approach. There are different map projection techniques, that either preserve angle, area, or length satisfying different properties and usages in cartography. These techniques flatten the earth into a 2D image, hence, abstract along the spatiality axis. Researchers have also applied common map projections on molecular data [35]. Their technique deforms the molecular surface into a sphere which is then projected to a 2D map. Metro maps are abstracting spatial characteristics to a degree that the topological properties are highlighted. The tracks of the metro system are at different depths. The depth is a detail that is completely irrelevant to the traveler and is thus removed. Metro map layouting techniques [69] reposition the stations in a certain layout (e.g., using octilinearity) to increase interpretability. In medical volumes, spatial properties can occlude important parts. Anatomical flattening [33] approaches reduce the spatiality based on anatomical features, but still depict the full visual details.

2.3.3 Perceptual Correspondence through Controlled Transition

At this point, multiple representations are defined along the two abstraction axes. As the representation depicts the same underlying data, the user could establish the correspon-




Figure 2.6: The geometric abstraction of the placenta along the spatiality axis. The structure is flattened to gain a standardized view on the fetal and the maternal sides. In the right end of the axis, the textual description contains little spatial information.

dence cognitively, but unnecessary cognitive effort would be the undesired consequence. As the transition changes the geometry of the representation, it is also referred to as geometric transition. In PlacentaMaps (**Paper A**), it was feasible to use CMVs to depict the small number of views. Unfortunately, for DNA nanostructures, the required screen space would be too large to fit all views. Hence, CMVs are not applicable in that case study. To depict the correspondence, a continuum between the representations has to be established. This continuum can depict the relationships between the representations explicitly. Using animation the user can observe a continuum between the isolated representations. This way, the user can perceive how one representation transitions into another one. Traditional animation relies on staging or on a step-wise method to guide the attention of the user to certain events. However, the constant speed of such an animation might not necessarily match the pace under which the user can perceive and process the geometric transition. In **Paper B** and **Paper C**, a user interaction element is provided to allow users to directly manipulate the animation step by step. For this control, the user is required to dedicate time to this interaction task. However, this control over the transition is not only necessary to depict the correspondence between views, but it enables the user to stop at the intermediate representations and perform modeling and analysis tasks. The user is not required to use the transition and could also jump directly to another representation. Hence, this approach would not cause additional interaction efforts in comparison to selecting a representation from a list.

Transition along the visual detail axis: This transition depicts how different visual detail levels are related to each other. As higher visual detail levels depict aggregated values, the transition animates how elements in one level are composed of elements of another level. For example, as one goes from high visual details at the atomic scales towards low details of entire structural entities, the transitions of the geometric primitives are achieved by interpolating the shape parameters and colors. Figure 2.7 shows the superimposed intermediate states between the atomic and the nucleotide representations. The defined atomic and nucleotide representations can be regarded as keyframes and intermediate frames are created by interpolation. The atoms are the constituents of nucleotides, which can be depicted by this controlled transition. As one can observe in Figure 2.7, there are more visual objects at the atomic than at the nucleotide representation. The relationship between one abstraction level to another is thereby either 1:1 or 1:n. In this case, one nucleotide consists of several atoms. The intermediate representation is thereby created by interpolating the parameters of several atoms to the same nucleotide. The transition thereby behaves as if several visual objects morph into one, going from high to low visual details. Figure 2.8 depicts the transition between single strand and double strand representations in ten steps. Another type of transition is blending between representations. The advantage of this method is that both representations are shown at different transparency levels. This method could be used to indicate that high-level structures are comprised of lower level structures. However, this approach might create visual clutter, since representations are blended with each other. Also, the correspondence between the representations is difficult to observe using this approach. Hence, the transition through interpolation of shape parameters and colors was chosen that creates a morphing animation.

Transition along the spatiality axis: The correspondence between different spatiality levels is more difficult to mentally establish if all spatial positions of the visual objects are different. The transition is achieved by interpolating the positions of the visual objects between spatiality levels.

In **Paper A**, the flattening of the placenta introduces little distortion by splitting first into the maternal and fetal side. Due to the low distortion (shown in Figure 2.6), the correspondences can be easily established through the CMVs. The controlled transition will not necessarily improve the perception of correspondence in this case. This highlights the case when transitions are not needed. In addition, when comparing multiple placentas, the correspondence of features cannot be established, which makes transitions infeasible. Therefore, a juxtaposed view was chosen here.

However, the 2D schematics in **Paper C** require deep knowledge of the inner workings of DNA in the context of nanotechnology.

Here, intermediate steps have to be made explicit when abstraction is performed along the spatiality axis. The animation along this axis depicts correspondences through the transition between two spatially abstracted representations. During the transition, only spatial locations of the elements are changed. An important goal for such a transition is to keep the distances between the locations as small as possible. The less the visual




Figure 2.7: Superposition of ten steps of the transition from atomic to nucleotide representations. The transition shows how the atoms morph together to form the beads representing the nucleotide. The abstraction goes along the visual detail axis.

elements are translated and rotated, the better the overall animation can be perceived. Figure 2.11 depicts an intermediate state in the transition from a 3D spatial structure to a 2D schematic representation.

2.3.4 Continuous Abstraction Space

Various features are abstracted into representations that go from concrete to abstract. The seamless transition and interpolation of the defined representations make this abstraction space continuous. Picking any location in this space outputs a representation. The spatiality and visual detail axes are abstracting different values. On the spatiality axis, the transitional variable controls the interpolation of the positions of the visual objects. On the visual detail axis, the transitional variable controls the interpolation of the shape parameters and colors. These two axes are used to construct a coordinate system where at discrete steps, representations are defined as shown in Figure 2.9.

Depending on the number of details removed, linear order of the representations on one axis can be established. For the visual details, this linear order naturally emerges from the hierarchical organization of the data. For example, in DNA, the suggested order is atoms, nucleotides, single strands and double strands. For the spatiality, it depends on the number of used dimensions, going from 3D, 2D, to 1D. A challenge arises when within one level of detail, several representations exist. For example, in **Paper B** the several representations for the atoms and bonds are placed along the visual detail axis.

The order is hereby at these local parts not along the direction of reduced complexity, as these representations depict different features but not necessarily fewer features. Along the spatiality axis, the same case would be if there are several 2D layouts. The natural hierarchy of the data would suggest only a global order, but at local parts, another strategy has to be employed to order. For example, in **Paper B** level 0 and 1 depict different features and neither representation is necessarily higher or lower in complexity. In that case, the strategy was to let the size of the visual objects decide. Smaller visual objects are ordered lower than larger ones, as this is anyway the trend along the entire abstraction axis.

One could think of a more modular approach in the future, where representation can be plugged into an abstraction space and at given levels of detail, certain representations can be chosen. This would provide the user with the possibility to customize and design his or her own abstraction space. So far, the distance between two neighboring representations is uniform. In case the relationship between two neighboring representations is more complex and the transition cannot be perceived in the given distance, smaller step sizes can be made using one spinbox for each abstraction axis. In terms of interaction, one disadvantage of using the spinboxes for transition is that only one variable is changed at a time and thus the movement is only parallel to the axes. In the future, one could imagine a non-linear, perceptually uniform abstraction space that is taking the number of changes into account. The animation would make smaller steps in those transitions that are harder to perceive, for example, if many visual objects are changed in appearance or in location.

Figure 1.3 shows the continuous abstraction space and the part that is covered by the resulting approaches of this thesis. This figure also depicts the progression of the idea. It started first with representations in **Paper A** that can be placed at 3D and 2D in the spatiality axis, but there were yet no transitions that show the intermediate representations. With the second paper (**Paper B**) a continuous abstraction approach of the 3D structure was developed. This approach covers the visual detail axis of 3D representations. In the last paper (**Paper C**), the DimSUM approach provides means to move freely in this abstraction space and adjust the abstraction level on the spatiality and visual detail level accordingly.

2.3.5 Controlling Abstraction

Using animation that depicts the transition from one to another representation makes the establishment of correspondence directly perceivable. The control of the abstraction level is important as it allows the user to adjust the pace of the animation to his/her cognitive ability to understand the transition. In contrast, an uncontrolled transition would be an animation where the user cannot control intermediate steps. In **Paper B**, the control over the animation is given to the user through a slider (Figure 2.8). The user can move the bar on the slider to adjust the representation in real-time. A slider is a user interaction widget that controls one position on an axis, which corresponds to one of the abstraction axes, as described in Section 2.3.2. The slider also encodes a direction,

for example from concrete (left) to abstract (right). Using one slider per axis is not ideal if there are two abstraction axes. The issue here is that to change the representation, the values have to be set on two sliders. Instead, the approach in **Paper C** directly changes the abstraction level along the two axes by displacing a point in a 2D panel to control the abstraction as shown in Figure 2.9. One current drawback of this approach is that the user has to divert the attention from the representation to the abstraction space. The user has to move the mouse in the user interface element, while the representation itself changes with every movement. A very interesting extension of this user interface could be marker menus [36]. A context-dependent menu that appears at the selected location on the object to control its representation would enable the user to rest his or her attention at the position. So far, the abstraction space is planned to contain all abstraction levels and the user can decide to use a subset of them for a given task in mind.




Figure 2.8: A slider to control the transition variable for one abstraction axis. Representations with decreasing visual details (atoms, nucleotides, single strands, double strands) are defined on the slider that represents an abstraction axis. At the top, the transition between the single strand and the double strand representation in ten steps is shown.

2.3.6 Guiding the User through the Abstraction Space

The proposed abstraction space contains many representations that are applicable for different tasks. For some tasks, only a subset of representations might suffice. For example, there are tasks, such as creating double-stranded DNA, that require high-level representations with low visual details. Other tasks, such as modifying atoms, require representations with high visual details. The user would need to first explore and discover locations (representations) that can depict the features of interest. An indication of relevant representations for specific tasks could guide the user's attention towards them.




Figure 2.9: The abstraction space as an interaction panel consists of the visual detail and the spatiality abstraction axis. Changing the position of the icon controls the transition in the abstraction space in real-time.

Similar to a sightseeing trip, there are only a few interesting locations and a guide could recommend them to the user. Recording the interaction with the abstraction space by an expert user, insight into the behavior can be acquired. The advantage of this approach is two-fold. First, understanding the behavior of a user, especially an expert user, can reveal insights that can be used to further improve the visualization system. Another, maybe even more important idea, is the transferability of a user’s experience. A novice user can learn from a more experienced scientist, where the relevant representations can be found. A visualization of the interesting locations could immediately guide the attention of a user to certain areas that others have found relevant.

Potentially interesting locations could be indicated by the amount of time a user spends in viewing a particular representation, as done in **Paper C**. This metric might be not the best indication of *interest* for one user, but in a setting where multiple users investigate the dataset and perform the same task, it could be more expressive. What is certainly necessary is to consider the context of the recorded interactions. Providing the context information would be crucial for a meaningful interpretation of the proposed recorded maps. For example, for the same dataset and the same task, a novice user would behave differently and spent time at very different locations in the abstraction map than an expert user. In **Paper C**, users were familiar with the dataset and the representations when their interactions were recorded. In an exploratory scenario, there would be also time spent at locations that are less relevant. In the future, one might think of additional metrics, such as time-stamps to record a full history of the interactions. Also, it would be very interesting to give the user the possibility to label the recorded interactions

retrospectively.

As any location in the abstraction space is associated with a representation, the system could record the time an expert spent viewing. The gathered data can then be loaded and used for augmenting the abstraction space. The values can then be regarded as a scalar field. Figure 2.10 shows how the abstraction space can be augmented with such information. The isolines visualize the scalar field and indicate the interesting locations like in a height map. Using this metaphor, the peaks are the most interesting locations with the representations of interest. Besides, particular interesting locations can be saved and depicted by icon-based visualizations as shown by the augmented images in Figure 2.10. The arrows indicate which path through this large abstraction space can be chosen for a particular task. In **Paper C**, an augmented abstraction space was implemented. The approach is tested by recording experts while they analyzed different nanostructures using the proposed abstraction space. Surprising insights could be derived from the analysis of the results.

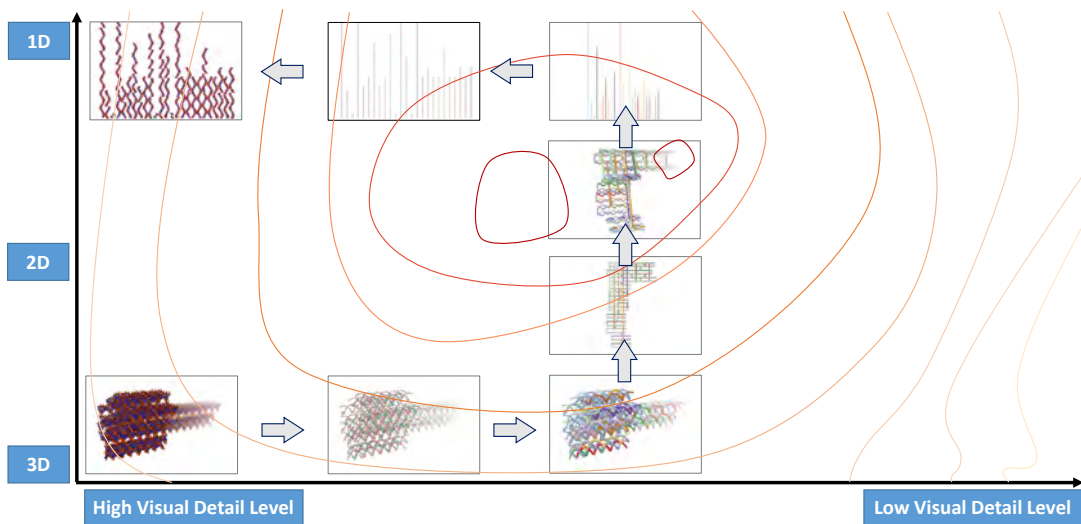


Figure 2.10: A concept sketch of an augmented abstraction space that contains many representations. It is augmented with additional information as shown here. The isolines depict a height map concerning interesting locations, while the icons depict an explicit path that a user took for carrying out a particular task.

2.3.7 The Usefulness of the Intermediate Representations

The recorded user interactions with the abstraction can also be incorporated as part of the evaluation of the proposed approach. For example, the idea of using intermediate representations as means of transition and creating correspondence between representation

can be evaluated by recording the experts. The recorded could easily depict, whether the expert spent time at the intermediate representations or not. By evaluating the interactions of experts with the abstraction space while they carried out analysis tasks, interesting insights into the preferences of the domain scientists for the intermediate representations could be gained. The intermediate steps are meant to depict the transformation between two representations. However, the data shows that certain intermediate representations are viewed for an extended time. The domain experts reported that the intermediate position would also provide additional insight into the data. The users appreciated the intermediate positions as they revealed a combination of two neighboring representations in certain cases as explained in **Paper C**. The insights could be made along both abstraction axes, as described in the following.

Abstraction along the visual detail axis: In Figure 2.12b the transitional state between the nucleotide (Figure 2.12a) and the single strand representation (Figure 2.12c) is shown. While Figure 2.12a shows the nucleotides through the individual spheres, Figure 2.12c depicts the DNA strands through the differently colored tubes. The transitional state enables the user to see the individual nucleotides as well as their affiliation to the single strands as indicated by the coloring.

Abstraction along spatiality axis: Figure 2.11b depicts the transitional state between a 3D structure and a 2D schematic representation. While the 3D structure in Figure 2.12a depicts the helicity of the DNA, the schematic view in Figure 2.11c only shows the connectivity through the crossovers. According to the domain expert, the transitional state in Figure 2.12b enables a view on the preserved schematic 2D layout, while already being able to observe structural properties, such as the helicity.

These two examples indicate how useful it is to quantify user behavior and subsequently to improve upon existing visual encodings. These insights would have been hidden from us, without the quantification of the experts' interaction and the augmentation of the abstraction space. In the future, the recorded data could provide more than a scalar field from user sessions. A vector field indicating the actual paths taken by experts would provide further interesting insights about the sequence of representations that are used for a task.

2.4 Integrated Modeling through Geometric Abstractions

A key purpose of the proposed visualizations is to enable the user to interactively model structures.

There exist commercial and open-source computer-aided design tools [5, 8] for 3D modeling with sophisticated interaction techniques. On first thought, one could assume that they might be sufficient for the interactive modeling nanostructures. These CAD tools already provide many 3D interaction and modeling techniques that would be useful for the design of nanostructures. However, the problem is that models created with these software do not carry the semantics of DNA molecules. The models would be just some geometries.

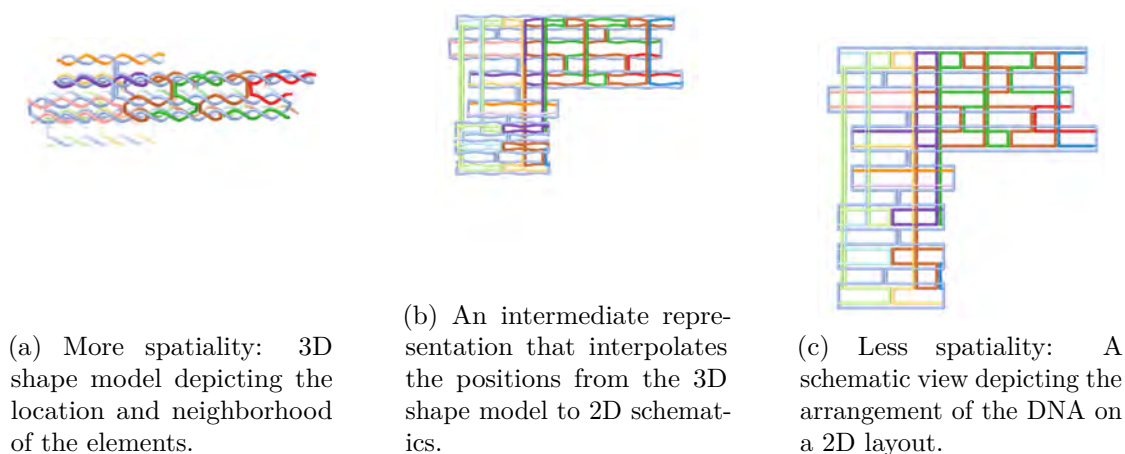


Figure 2.11: The transition of a DNA nanotube along the spatiality axis from the 3D shape to the 2D schematics. (b) depicts an intermediate representation. It shows the correspondences and provides additionally a combination of the 2D schematics and the 3D shape.




Figure 2.12: The transition of a DNA nanotube along the visual detail axis from the nucleotide to the single strand representation. The intermediate image still depicts the nucleotides as spheres, but the transition towards the single strand color already provides an indication of the nucleotides' affiliation to the corresponding single strands.

For example, these semantics are needed to compute and visualize stability measures of the nanostructures and to aid the modeling operations, based on the structural properties of the DNA. In the case of this thesis, a DNA molecule can be represented by a tube, a chain of beads, a collection of spheres, depending on the chosen abstraction level, as described in **Paper B** and **Paper C**. The DNA nanostructures, although just several tens of nanometers large, still consist of a great number of atoms and features. On that level of detail modeling operations would be time-consuming and tedious. An early insight was that there is a big difference between visualization for analysis and for modeling. The geometry used to present the features have to be visualized in a way that is suitable for being interacted with. Fortunately, the hierarchy of biological structures enables the

application of modeling operations on groups of elements. The selection and picking of elements of a group alone is an elementary task in modeling. Besides the size of the structure, the high density of molecular data is also a challenge that needs to be solved through visualization. Another problem is that the elements hide others. With the proposed method, the user is equipped with an integrated solution to navigate to and apply operations on any representation of desired visual detail and spatiality, as shown in Figure 2.1 and Figure 2.3.

2.4.1 Abstraction-based Selection


The selection of elements in the scene is the very basic interaction that needs to be enabled before modeling operations can be applied. For analysis tasks, the selection and the subsequent highlighting of the selected elements also enable the user to see correspondences across different representations. For modeling, the task usually involves selecting one or more elements from the scene. Selecting multiple associated visual elements can be challenging if there are many of them and the scene is cluttered. However, using the abstract representations, the selection can be done in different visual detail levels (Figure 2.13) or spatial dimensions (Figure 2.14).

Selection along the visual detail axis: The abstraction process creates representations with different levels of visual detail. As a consequence, high-level features and their representation using fewer geometric primitives can be facilitated to selected lower-level features. Furthermore, primitives, such as spheres and cylinders have great advantages if the elements have to be selected, translated, and rotated. For example, when visualizing nucleic acids, the traditional ribbons and sticks (Figure 2.13a) are better suited for depicting the orientation and position of nucleotides, but the selection can be better performed on the spheres or cylinders (Figure 2.13b). Using spheres, it is easier to distinguish one nucleotide from another one. In Figure 2.13c entire strands are represented by a simple tube that is easily distinguishable. High abstraction levels are even more helpful, such as in Figure 2.13d, since they enable the user to pick a base-pair or the entire double strand.


Selection along the spatiality axis: The advantage of spatiality levels are the different examination possibilities of the visual objects. Each spatiality level can reveal certain features of the data better than another level. The advantage that makes the inspection process more effective, also makes the selection easier. In Figure 2.14, one DNA strand is selected and highlighted across different spatiality levels of DNA nanostructures. 1D (Figure 2.14a) is particularly advantageous to select a strand of a certain length or with a certain sequence, which can be particularly well observed in the linear arrangement. 2D (Figure 2.14b) shows how the single strands cross over from one double strand to another. 3D (Figure 2.14c) depicts the targeted architecture and enables one to easily access the spatial location of the strand. Selection can be performed at the spatiality level that is optimal for choosing the visual element with the appropriate set of characteristics.




(a) Ribbons and sticks as nucleotides [27].



(b) Spheres and cylinders as nucleotides.



(c) Intertwined helical tubes as DNA single strands.



(d) Spheres as base pairs and DNA double strands.

Figure 2.13: (a) shows a common ribbon visualization of nucleic acids, which is not optimal for interactive modeling of nanostructures. (b-d) show geometric primitives, such as spheres and cylinders for efficient selection of grouped atoms. In (b-d) the selection of atoms can be facilitated on a higher visual detail level.

2.4.2 Abstraction-adaptive Modeling

The goal of modeling is to modify, remove, or add elements to an existing structure or create entirely new ones. The benefit of the proposed method for the user is that modeling tasks can be performed at an ideal abstraction level where the feature of interest can be observed the best. To exemplify this, a use case is shown in Figure 2.15. In Figure 2.15a a DNA nanoshell is depicted. A double-stranded cross-bracing should be added at the top part of the shell. Performing this task on the atomistic representation would be tedious, as there are many atoms to insert. The atomistic representation would not be

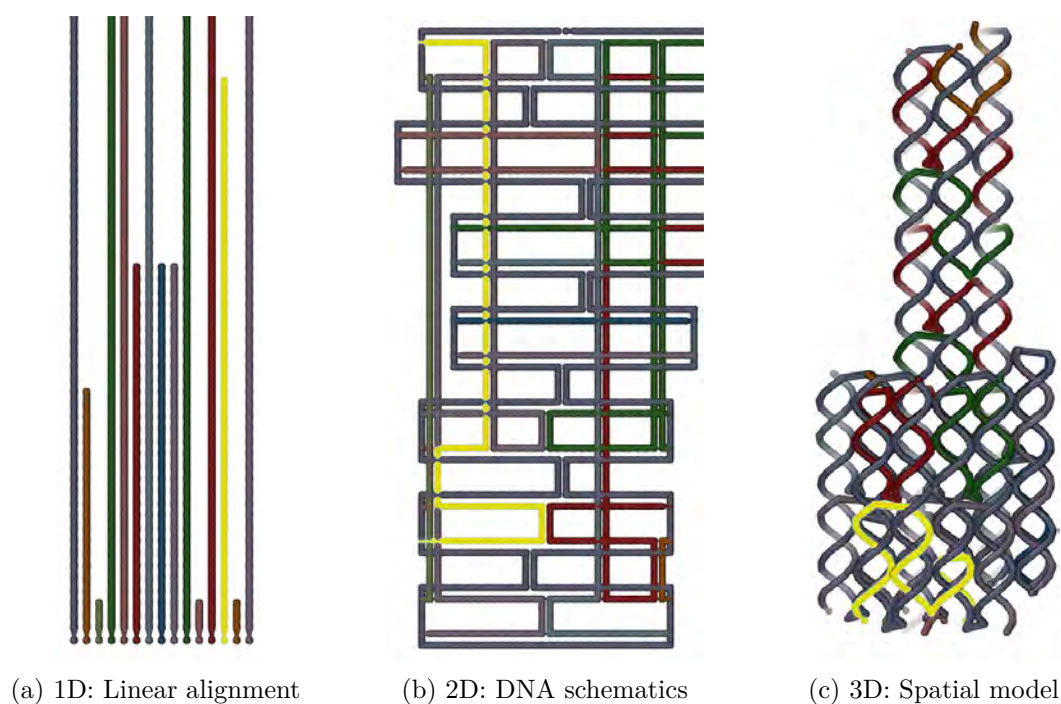



Figure 2.14: The selection across different spatiality levels. The same strand can be selected in any of the 1D, 2D, or 3D representations.


appropriate as the visual detail level is unnecessarily high for such a task. Instead, the double strand at a low visual detail level (Figure 2.15b) provides a clear representation of the width and length of the DNA double helix. This task does not require any inspection of the atomistic details. Therefore, the user can choose to operate at a high-level of abstraction, as shown in Figure 2.15c. The user does not have to worry about how the modeling operations will change automatically other levels of abstraction, as the operation automatically propagates across the entire abstraction space, as demonstrated in **Paper B**. The advantage of this abstraction-adaptive modeling is that the user can stay in the abstraction level that is most appropriate for a given task.

2.5 Evaluation

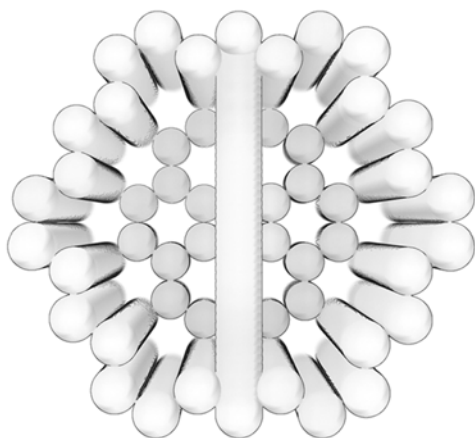
The research carried out for this thesis is application-driven and, hence, the targeted users were involved in the evaluation of the proposed approaches. For the two case-studies, different evaluation strategies were chosen. In the PlacentaMaps project, a two-part survey was carried out with domain scientists with various backgrounds, e.g., in fetal/placental imaging, gynecology, clinical data analysis. Their common interest is in the clinical research of the placenta. In the first part, they rated three aspects (structure-aware slicing, comparability, flattening) of the proposed technique based on pre-recorded videos. This was done remotely. The second interactive part was conducted by providing




(a) Atomistic representation before a cross-bracing is inserted.



(b) Inserting the abstracted cross-bracing on the atomistic representation.



(c) Using the same level of abstraction for existing and inserted structure.



(d) The atomistic details are automatically generated based on a high-level abstraction.

Figure 2.15: Modifying a nanoshell by adding a DNA double strand as a vertical cross-bracing. (a) the structure consists of a large number of atoms, resulting in high visual details. (b) the double strand is interactively inserted, using an abstracted representation where the visual detail is low. (c) shows that insertion can be performed in other visual detail levels. (d) depicts the propagation of the insertion to other representations, in this case to the highest visual detail level.

scientists with the tool and having them test the functionalities. Through semi-structured interviews, the feedback was collected, which is described in **Paper A**. Overall ratings of readability, clinical applicability, research applicability were also discussed with the experts. As this approach flattens the placenta, the area distortions were also measured.

There are some lessons learned from this approach. The goal of this project was

straightforward and the evaluation was a matter of how satisfied the experts are with the proposed technique with respect to their own tasks. For this, input from the respective scientists was collected. The requirements were defined by high-level tasks [46], such as "I want to understand the attachment to uterine wall", and by low-level tasks, such as "I want to count and compare the cotyledons on the maternal side of the placenta". The feedback for the prototypes was acquired during milestones, which could be a month or longer. Sometimes plenty of implementation efforts have already been invested until it was realized that the chosen prototype will not sufficiently solve the requirements. In one instance, a research prototype was developed, where the placenta volume was virtually inflated into a sphere and then using an azimuthal map projection approach placenta was flattened into one 2D image. Although the result was a 2D image, the introduced distortion was simply too high, making the visualization design not sufficient. This mistake could have potentially been avoided if either more regular meetings or extensive regular surveys would have been conducted. However, extensive surveys require effort and time from both sides, the visualization researcher and the domain experts. It is time-consuming to create the survey and to evaluate it, and it is organizationally challenging to involve several experts, especially in an international and multidisciplinary collaboration. A discussion of potential pitfalls is provided by Sedlmair et al. [60].

The change of project context from PlacentaMaps to Mara provided a chance to establish a new type of collaboration and also evaluation strategy. In this project, the setup enabled me to be directly embedded in the team of domain scientists. While it was clear what the overall goal of the entire project is, i.e., "to create a visual tool that enables experts to design a catalytic nanorobot", the actual visualization and interaction-related problems had to be derived first. These tasks are described in **Paper B** and **Paper C**. A particular challenge was that the nanotechnology domain did not receive much attention from the visualization field and therefore, there was little research work to start with. Also, the domain itself is young and quite complex. Much domain knowledge needed to be learned to understand the underlying problems. Fortunately, the embedded setup enabled an extremely fast feedback cycle that proved to be very beneficial for several aspects. The domain language could be learned quickly, which increased the understanding of the domain problems. Getting feedback was as simple as asking a colleague in the shared office, or accompanying and observing the workflow of the in silico designers or the researchers working in the wet lab. Required data was provided instantly. This led to a rapid prototyping stage, where ideas were quickly gathered and the utility of the visualization concept could be quickly tested on the available testers. Another unforeseen consequence of such an embedded collaboration was that both sides would understand the other better. It was not only to me, the visualization researcher to get an increasingly in-depth understanding of the challenges of the domain, but the domain expert would also become literate in the visualization field. In contrast to the PlacentaMaps project, this collaborative setup avoided wrong assumptions, which then saved valuable implementation time.

However, such an embedded setup does not come without challenges. The likely biggest

challenge are diverging goals. Pursuing a PhD in visualization, one aims for publications in this field. However, there is the goal of the project, where a tool has to be developed that supports the domain scientists in their work. There is much engineering effort needed, that could potentially stand in the way to publishing visualization research. This needs to be communicated early on with the principal investigator of the domain so that one does not become a technician for the domain experts. The work arrangement was challenging, but the immediateness of the feedback on the usefulness by regular users was essential for the progress in the research.

The visual tool was then deployed to in-house scientists who provided regular feedback and input. In **Paper B**, the proposed multiscale approach was evaluated in a use case where a tremendous increase in time-efficiency could be achieved when comparing to the previous method. The use case was the connecting of components into higher-order structures. As the existing and novel methods are different, it was also important to show that the produced result of the proposed approach is valid. This was done by comparing the produced sequences, which are the final outputs of the existing and novel method. The sequences are sent to wet-lab experiments in the next step. After the experts carried out the self-assembly, transmission electron microscopy images are acquired to show the fabricated structure and demonstrate the design, using the proposed visualization concept was successful. In **Paper C**, two use cases were demonstrated. For example, the adding of bridging strands requires domain scientists to operate in 1D, 2D, and 3D. Here, the recording of the interactions in the abstraction space was providing quantitative results that are used for further questioning.

Reflecting Remarks

What I want to talk about is the problem of manipulating and controlling things on a small scale.

Richard Feynman

I described geometric abstraction as an approach to create effective representations and to integrate them in a conceptual abstraction space. The idea of abstraction of the spatiality was developed from the flattening of the placenta. Later, this idea was developed further when the application domain changed into nanotechnology. The resulting visualization system is not only a contribution to the visualization field, but also to DNA nanotechnology. The visualization system is an integral part of the Adenita toolkit [40]. This toolkit is realized as an extension to the SAMSON framework [47] that provides additional modeling features for proteins combining with the structures created with Adenita. As a co-developer, I created the visualizations and their features in tight collaboration with Adenita users in the course of two and a half years. It was crucial that I became embedded into the MARA team, supporting the scientists with visualization and interaction tools. Under these circumstances, they applied the visualization concepts regularly, as they use Adenita for algorithmic design as well as for modeling of DNA nanostructures with novel shapes and functions, such as a DNA nanorobot.

3.1 Impact

Medical visualization, the context of the first use case, has been an important part of the scientific visualization field. Concepts and tools developed here can now be applied to emerging fields, such as nanotechnology, as demonstrated in this thesis. DNA

nanotechnology has potential applications in medicine and could fundamentally change the way how drugs are designed and administered in the future. First drug delivering robots have been already created [15, 1]. The aim is to model arbitrary functional structures at the nanoscale. The design is a highly computational process, as the scientists need to create appropriate models in silico. Visualization is thus fundamental in inspecting structural characteristics and applying interactions for subsequent modeling tasks. As discussed in **Paper B**, the Adenita toolkit goes beyond the concept stage, and it is applied to design nanostructures that are fabricated in laboratory experiments by the collaborators. Also, the involved scientists used the visual tools and developed several novel nanostructures that could not be designed with the existing ones. These structures are now under experimental validation. The development progress of Adenita was presented regularly at workshops, where it attracted the interest of other research groups that focus on DNA nanostructures. In the meantime, the alpha version of Adenita has been deployed to several research groups for helping them in their research. A beta release of Adenita to the public will be available soon [4].

3.2 Discussion

The use cases investigated here provided a rich spectrum of challenges that geometric abstraction tackles. The method I proposed puts the cognitive efforts into focus and considers mental models in the design of novel visualization systems. I discuss the different steps in the geometric abstraction in the following. Considering the mental models could be useful in collaborations, when there are established existing workflows. One should still keep in mind that sometimes, the strategy of keeping the cognitive effort low might not be the best one. There are cases where existing workflows have to be questioned and radically new representations have to be designed. They might be unfamiliar first and require a strong effort in learning. If feasible, the transition that depicts the transformation from the familiar representation into the novel visualization could ease the steep learning curve for the latter.

The general application of my proposed approach to facilitate various abstraction levels for depicting different features in complex data depends on several aspects. First of all, the complexity of the data should make it necessary to create multiple representations. Secondly, for the abstraction of visual details, the data needs to have a certain hierarchical organization. When abstracting its spatial composition, different levels should also exist for the data and clear transitions between the different representations need to be defined, need to be meaningful, and the resulting representations need to be managed. This may not always be feasible in all cases, and alternative approaches may be easier to realize. For example, a recent approach [24] used *visual embedding* of one scale representation in another for showing the scale hierarchy from the nucleus of a cell to the atomistic composition of the DNA, without the need to define specific geometric transitions, as described in Section 2.3.3. On the other hand, dedicated geometric transitions may allow viewers to better understand how elements in one level related to the respective elements in the next.

Proteins are large complex molecules that consist of amino acid chains. Their composition is similar to DNA and can be described by the primary, secondary, tertiary and quaternary structure. The similarity invites the application of geometric abstraction, as described in this work. Furthermore, with the rise of protein origami [39], the requirements between these two domains become very similar.

Many application scenarios of molecular data were provided already. To view this approach in a broader context, one could also apply the geometric abstraction on astronomical data. Generalizing the idea of the guiding abstraction space, it would be very interesting to stir the attention of the user to interesting appearances in the vast outer space. Astronomical systems distinguish from molecular data by the sparseness of the outer space and the large size ratio between celestial bodies. Furthermore, the geometric transition would be challenging to realize due to the large distances between celestial bodies.

The concepts described in this work could be applied to various other domains where the above-mentioned aspects are fulfilled. A metabolic pathway visualization [55, 29] depicts a large graph of metabolic reactions. The complexity is due to a large number of nodes and their reactions with each other through the metabolites. Many reoccurring high-level motifs are studied by experts. Such a network could benefit from several abstraction levels and the geometric abstraction to analyze these networks.

One could also question whether the transitions should always be depicted. First, the control itself enables the user to adjust the animation step by step. Perceiving the correspondence might not always be necessary. From observations of the workflows, the user is more likely to skip the transition when he or she is already more familiar with the dataset.

The transition is based on the interpolation of the geometric primitives that are used by the representations. Cylinders and spheres are often used in molecular visualization. The geometric transition might be challenging to realize for other types of primitives. For example, the transition between a sphere and a cube might be more challenging. In such cases, blending could be another type of transition that might be applied as it does not adjust shape primitives for interpolation. The current approach allows the user to pick one location in the abstraction space and the system adapts the representation accordingly. This is a limitation when features of different abstraction levels have to be compared, as only one representation is depicted at a time. However, complementing the proposed approach with CMVs could handle such tasks.

When changing the abstraction along the spatiality axis, the transition can be hectic if many objects change spatial location at the same time. Although a brushing tool (see **Paper C**) is provided, a sequential method could improve the perception of the animation.

I described the abstraction of visual detail and spatiality as the two dominant ways, how spatial structures can be changed in their geometry and visualized. I demonstrated the idea in use cases from two different fields. Geometric abstraction only concerns the model in object space. A generalization of this idea is to extend the concept of

geometric abstraction to a general abstraction strategy that considers many other viewing parameters. After choosing the appropriate abstraction level, an expert user is also navigating the camera to inspect interesting features. The control of the camera could also provide a novice user with an indication of interesting locations in the abstraction space. The camera movement by the user can be added as another axis in the abstraction space. The guidance in the abstraction space could be further extended by indicating the relevant representations for a given task, as well as navigating the camera to an advantageous position where important features are located. Managing the visibility is another way, how the complexity can be reduced. In this work, the clutter is dealt with abstraction. However, a combination with advanced visibility management techniques, as already indicated in **Paper B**, could provide the user with even more effective design. For example, one could imagine an abstraction-adaptive visibility management technique, that changes the visibility of visual objects depending on the abstraction level. Creating different visual encodings for the same data at different abstraction levels and the creation of meaningful seamless transitions could be regarded as tedious and time-consuming. Moreover, for different tasks within one domain as well as for applications from different domains, the specific sequence of abstraction levels may not always be the same as it very much depends on the mental models of the users. This means that also different transitions may need to be defined depending on the given abstraction level sequence. The advantage, however, is that each representation can be optimized in conveying only certain features. Moreover, the integration into an abstraction axis diminishes the additional interaction efforts to switch representations, as the abstraction-space panel enables users to quickly transition between visualizations. The process of guiding viewers towards interesting locations in the abstraction space as I have implemented is currently limited by acquired data. The meaningfulness of time spent at locations depends on the context of the interaction and has to be carefully interpreted when applied as an indication for interesting locations. The strategy of geometric abstraction is driven by the application field. In this way, the effectiveness is also demonstrated by the long-term users that test and apply the resulting software. It also shows that the representations are carefully designed and integrated to suit the tasks of these experts. One might assume that only a few expert users exist. DNA nanotechnology is certainly a very specialized field, however, the therapeutic applications are very promising as previously stated. Furthermore, an increasing number of research groups started to apply the technology. At the time of writing of the thesis, already more than 20 scientists requested access to the current prototype of the developed software toolkit. In addition to the in house scientists, an alpha version of Adenita was deployed to several research groups in Europe, supporting them in the design of novel nanostructures. The potential impact of the technology itself and the growing need for computational tools makes the research on the visualization for this field timely and relevant.

3.3 Conclusion and Future Work

The results suggest that complex scientific data, such as molecular structures and medical volumes, can be abstracted into simplified representations when their level of visual details and spatiality are reduced. The order, given by the amount of features removed, is employed to create an abstraction axis of the conceptual abstraction space. The resulting visualization and interaction system enable scientists to study and design DNA nanostructures in a way that was not possible before. Although a prototypical toolkit has already been developed in the course of this thesis, the research will continue. Figure 3.1 provides a montage of the visualizations that can be achieved with the proposed method. The figure also sketches the future direction. Depicting the different levels of abstraction locally in a focus region without changing the context could improve the modeling process.

Geometric abstraction extends the literature in several ways. The continuous abstraction level and the creation of an controllable abstraction space shows great potential in how we can deal with multiple representations. The controlled transitions are achieved by interpolating the shape parameters, colors, and positions of the geometric primitives. In scenes with many visual details, the transition might be cluttered as many elements change at the same time. This problem is especially prevalent in transitions along the spatiality axis. As Fisher [18] notes, it might be better to spare the users the animation if everything on the screen is moving around. There is still definitely need for more sophisticated transitions than the proposed interpolation-based method in this thesis. One could think of a local method that preserves the context, but transitions a focus region. The selection and highlighting of elements might mitigate this problem and make following their transitions easier. Another idea is to dissect the transition into multiple stages, as used in animated transitions between different types of information graphics [26]. For example, the transition from 2D to 1D in **Paper C** could be dissected into three stages, i.e., first the untwisting of the double helix, second the straightening of strands and finally, the alignment of the straightened strands in 1D.

The advantage of the abstraction space is to have many representations at hand and being able to easily switch to a representation that suits a task the best. The transitions ensure that each representation can be understood better by relating it to more familiar visualizations if there are any in the defined space. This idea can also be transferred to other domains. The idea of a controlled transition between representations is not necessarily restricted to medical and biological structures. In information graphics, different representations have advantages and disadvantages. Enabling the user to directly control the transition and observe the animation between two types of charts, could help users to better understand the relation between two types of visualizations, especially if they have different advantages and disadvantages. The control over the transition could combine the advantages of two visualizations. Even within the molecular domain, there might be data, where the proposed method can be applied. Although the behavior is fundamentally different, a protein's composition is similar to DNA. Instead of nucleic acids, the building blocks of proteins are amino acids. The modeling of protein in the context of DNA-protein hybrid structures have been already experimentally validated

[49]. So it might also be interesting to investigate how the proposed abstraction approach could be beneficial in this field. The same is true for RNA structures. RNA molecules are structurally even closer to DNA than proteins. Nanostructures built from RNA have also been proposed already [25, 20]. In conclusion, in the molecular domain alone, there are many potential applications of the proposed method.

A very commonly demanded feature by the users is the integration of molecular dynamics simulations into the visualization system. The simulation predicts whether the designed structure will self-assemble under the defined conditions. Such predictions could further improve the entire fabrication process itself.

One key result is that the different levels of abstractions act as the foundation for interactive modeling. With the result of this thesis, the first steps towards a comprehensive modeling pipeline are taken. In the future, the focus will be on the modeling process itself. It requires sophisticated visualization and interaction techniques to design the nanostructures with novel functionality. Helpful techniques, such as snapping and alignment tools, need to be implemented to make the design process more efficient. During interaction, helpful augmentations in the interface could create an enhanced and optimized modeling process. If invalid modeling operations are performed, the system could communicate this to the designer through appropriate visualizations.

One important feature of our modeling approach is the dependence of the representations to each other. Operations carried out in one representation will automatically impact others. The reason for this is the common underlying model. Right now, the interaction is applied by one user at a time. In the future, this approach could be extended to enable several users to operate in parallel and work on the same dataset. One user could work on a high-level architecture of a structure and another could focus on the detailed atomic modifications. This collaboration could be even further extended by sharing tasks across different modalities. One could imagine an expert working on a desktop computer using the 2D representations, while another is working in virtual reality using the 3D representations.

In a wider context, these ideas fit into the automation of the fabrication process of nanostructures. The *in silico* design is a crucial step and the visualization concept proposed in this work points clearly towards the digital fabrication of nanorobots.

3.4 Authorship statement

The articles listed in this thesis were published during the doctoral studies. They were developed in collaboration with multiple co-authors from the field of computer science, medicine, biology, and physics. These projects were interdisciplinary, the publications had contributions from domain scientists that helped in creating ideas, gathered data, provided feedback, and tested the developed methods. These contributions are reflected on the co-author list. However, the research was driven and carried out by the thesis

author, who realized the technical implementation and wrote the main parts of the publications that constitute this cumulative thesis.

- **PlacentaMaps project (Paper A)**

The exploratory nature of this project required the discussion and feedback of experts with various backgrounds. The domain expertise for this project was provided from collaborators in clinical data analysis, pediatrics, and gynecology from King’s College, London. For the volume segmentation of the placenta, we collaborated with researchers at the Imperial College, London. Furthermore, I collaborated with a leading group in medical visualization at the Otto-von-Guericke University Magdeburg, Germany. The idea was developed by the thesis author and later refined together with the co-authors Bernhard Kainz and Gabriel Mistelbauer. It was then evaluated with the involved team of medical doctors. They provided the data, segmentation, and background information, and gave feedback in the context of clinical research and diagnostics. The research and technical realization of the conceptual idea was done by the thesis author. Gabriel Mistelbauer supported the supervision and advised on the technical realization. Eduard Gröller led the project and directed the overall research.

- **MARA project (Paper B & C)**

I was part of MARA team at the Austrian Institute of Technology, which is specialized in applying DNA nanotechnology to develop molecular diagnostic tools. Here, I developed the ideas of both publications. With Inria (France), I collaborated on the topic of visual abstraction in the context of biology. The technical realization of the visualization concepts and the research was mainly done by me. Further responsibilities were the development of the software for the MARA project together with Elisa De Llano (a physicist), who helped with the implementation of the underlying data model for the visualization. The main supervisor Ivan Viola sketched the plan to carry out this research and guided the research vision. Co-supervisors Tobias Isenberg and Eduard Gröller refined subsequently the ideas and together with Ivan Viola helped to condense them into research papers. As this is a highly interdisciplinary project, the principal investigator Ivan Barišić as an expert in DNA-based diagnostics and modeling was the co-supervisor in the field of DNA nanotechnology. He and Elisa De Llano were also the main users of the developed methods. The other authors provided the data and helped to define the visualization concepts by regular discussions and by their feedback while working with the software.

The development of the thesis topic did not start and end with the publications in this cumulative thesis. The following papers were also published in the course of the PhD studies:

CoWRadar: Visual Quantification of the Circle of Willis in Stroke Patients.
Haichao Miao, Gabriel Mistelbauer, Christian Nasel, M. Eduard Gröller. *Proceedings of*

Eurographics Workshop on Visual Computing for Biology and Medicine (EG VCBM). Sep 14-15, 2015, 1–10. [Best Paper Honorable Mention at the EG VCBM 2015]

Visual Quantification of the Circle of Willis: An Automated Identification and Standardized Representation. [Haichao Miao](#), Gabriel Mistelbauer, Christian Nasel, M. Eduard Gröller. *Computer Graphics Forum*. 36(6): 393-404, 2016.

Multiscale Molecular Visualization. [Haichao Miao](#), Tobias Klein, David Kouril, Peter Mindek, Karsten Schatz, Meister Eduard Gröller, Barbora Kozlíková, Tobias Isenberg, Ivan Viola. *Journal of Molecular Biology*. ISSN 0022-2836, 2018, 1049-1070.

A Preview to Adenita: A Software Toolkit for the Visualization and Modeling of DNA Nanostructures. Elisa De Llano, [Haichao Miao](#), Tobias Isenberg, M. Eduard Gröller, Ivan Viola, and Ivan Barišić. *Poster shown at the 3rd Functional DNA Nanotechnology Workshop*. Rome, Italy. June 2018.

Interactive Visual Analysis for the Design of DNA Nanostructures. [Haichao Miao](#), Elisa De Llano, Ivan Viola, and Ivan Barišić. *Poster shown at the Nantech Workshop: Nucleic Acid Nanotechnology: From Algorithmic Design to Biochemical Applications*. Espoo, Finland. May 2019.




Figure 3.1: A nanorobot [11] is visualized with the proposed method. The figure comprises a summary of the representations and the possibilities for expert users.

Bibliography

- [1] E. S. Andersen, M. Dong, M. M. Nielsen, K. Jahn, R. Subramani, W. Mamdouh, M. M. Golas, B. Sander, H. Stark, C. L. P. Oliveira, J. S. Pedersen, V. Birkedal, F. Besenbacher, K. V. Gothelf, and J. Kjems. Self-assembly of a nanoscale DNA box with a controllable lid. *Nature*, 459(7243):73–76, 2009.
- [2] G. Andrienko and N. Andrienko. Coordinated multiple views: A critical view. In *Proceedings of the Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization*, CMV '07, pages 72–74, Washington, DC, USA, 2007. IEEE Computer Society.
- [3] F. Arias, L. Rodriguez, S. C. Rayne, and F. T. Kraus. Maternal placental vasculopathy and infection: Two distinct subgroups among patients with preterm labor and preterm ruptured membranes. *American Journal of Obstetrics and Gynecology*, 168(2):585–591, 1993.
- [4] Austrian Institute of Technology. MARA – Molecular analytical robotics assays. Website: <http://maraproject.eu/>, 2019. Visited in August 2019.
- [5] Autodesk. Autodesk Maya. Website: <https://www.autodesk.com/products/maya/overview>, 2019. Visited in August 2019.
- [6] K. Benirschke, G. J. Burton, and R. N. Baergen. *Pathology of the Human Placenta*. Springer Science & Business Media, 2012.
- [7] E. Benson, A. Mohammed, J. Gardell, S. Masich, E. Czeizler, P. Orponen, and B. Högborg. DNA rendering of polyhedral meshes at the nanoscale. *Nature*, 523(7561):441–444, 2015.
- [8] Blender Foundation. Blender. Website: <https://www.blender.org/>, 2019. Visited in August 2019.
- [9] R. Borgo, J. Kehler, D. H. S. Chung, E. Maguire, R. S. Laramée, H. Hauser, M. Ward, and M. Chen. Glyph-based Visualization: Foundations, Design Guidelines, Techniques and Applications. In *Eurographics 2013 - State of the Art Reports*, pages 39–63. The Eurographics Association, 2013.

- [10] A. Buja, J. A. McDonald, J. Michalak, and W. Stuetzle. Interactive data visualization using focusing and linking. In *Proceedings of Visualization*, pages 156–163, Los Alamitos, 1991. IEEE Computer Society.
- [11] C. E. Castro, F. Kilchherr, D. N. Kim, E. L. Shiao, T. Wauer, P. Wortmann, M. Bathe, and H. Dietz. A primer to scaffolded DNA origami. *Nature Methods*, 8(3):221–229, 2011.
- [12] G. Cipriano and M. Gleicher. Molecular surface abstraction. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1608–1615, 2007.
- [13] P. Decaudin. Cartoon looking rendering of 3D scenes. Research Report 2919, INRIA, 1996.
- [14] R. Desimone and J. Duncan. Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18:193–222, 1995.
- [15] S. M. Douglas, I. Bachelet, and G. M. Church. A logic-gated nanorobot for targeted transport of molecular payloads. *Science*, 335(6070):831–834, 2012.
- [16] S. M. Douglas, A. H. Marblestone, S. Teerapittayanon, A. Vazquez, G. M. Church, and W. M. Shih. Rapid prototyping of 3D DNA-origami shapes with caDNAno. *Nucleic Acids Research*, 37(15):5001–5006, 2009.
- [17] N. Elmqvist and P. Tsigas. A taxonomy of 3D occlusion management for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(5):1095–1109, 2008.
- [18] D. Fisher. *Animation for Visualization: Opportunities and Drawbacks*. O’Reilly Media, 2010.
- [19] J. W. Forrester. Counterintuitive behavior of social systems. *Theory and Decision*, 2(2):109–140, 1971.
- [20] C. Geary, P. W. K. Rothmund, and E. S. Andersen. A single-stranded architecture for cotranscriptional folding of RNA nanostructures. *Science*, 345(6198):799–804, 2014.
- [21] M. Glueck, K. Crane, S. Anderson, A. Rutnik, and A. Khan. Multiscale 3D reference visualization. In *Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games*, I3D ’09, pages 225–232, New York, NY, USA, 2009. ACM.
- [22] M. Glueck, A. Khan, and D. J. Wigdor. Dive in! enabling progressive loading for real-time navigation of data visualizations. In *Proceedings of Conference on Human Factors in Computing Systems*, 2014.
- [23] A. Gooch, B. Gooch, P. Shirley, and E. Cohen. A non-photorealistic lighting model for automatic technical illustration. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH ’98, pages 447–452, New York, NY, USA, 1998. ACM.

- [24] S. Halladjian, H. Miao, D. Kouřil, M. E. Gröller, I. Viola, and T. Isenberg. Scaletrotter: Illustrative visual travels across negative scales. *To appear in IEEE Transactions on Visualization and Computer Graphics*, 26(1), 2020.
- [25] D. Han, X. Qi, C. Myhrvold, B. Wang, M. Dai, S. Jiang, M. Bates, Y. Liu, B. An, F. Zhang, H. Yan, and P. Yin. Single-stranded DNA and RNA origami. *Science*, 358(6369), 2017.
- [26] J. Heer and G. Robertson. Animated transitions in statistical data graphics. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1240–1247, 2007.
- [27] W. Humphrey, A. Dalke, and K. Schulten. VMD: Visual molecular dynamics. *Journal of molecular graphics*, 14(1):33–38, 1996.
- [28] A. Kanitsar, D. Fleischmann, R. Wegenkittl, P. Felkel, and M. E. Gröller. CPR – Curved Planar Reformation. In *Proceedings of IEEE Visualization*, pages 37–44, 2002.
- [29] Kegg. Pathways. Website: <https://www.kegg.jp/>, 2019. Visited in August 2019.
- [30] S. Kmiecik, D. Gront, M. Kolinski, L. Wieteska, A. E. Dawid, and A. Kolinski. Coarse-grained protein models and their applications. *Chemical Reviews*, 116(14):7898–7936, 2016. PMID: 27333362.
- [31] B. Kozlíková, M. Krone, M. Falk, N. Lindow, M. Baaden, D. Baum, I. Viola, J. Parulek, and H.-C. Hege. Visualization of biomolecular structures: State of the art revisited. *Computer Graphics Forum*, 36(8):178–204, 2017.
- [32] B. Kozlíková, M. Krone, N. Lindow, M. Falk, M. Baaden, D. Baum, I. Viola, J. Parulek, and H.-C. Hege. Visualization of biomolecular structures: State of the art. In *Eurographics Conference on Visualization (EuroVis) - STARs*, pages 061–081. The Eurographics Association, 2015.
- [33] J. Kreiser, M. Meuschke, G. Mistelbauer, B. Preim, and T. Ropinski. A survey of flattening-based medical visualization techniques. *Computer Graphics Forum*, 37(3):597–624, 2018.
- [34] J. Kretschmer, G. Soza, C. Tietjen, M. Suehling, B. Preim, and M. Stamminger. ADR - Anatomy-Driven Reformation. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2496–2505, 2014.
- [35] M. Krone, F. Frieß, K. Scharnowski, G. Reina, S. Fademrecht, T. Kulschewski, J. Pleiss, and T. Ertl. Molecular surface maps. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):701–710, 2017.

- [36] G. Kurtenbach and W. Buxton. User learning and performance with marking menus. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '94, pages 258–264, New York, NY, USA, 1994. ACM.
- [37] G. Lakoff and M. Johnson. The metaphorical structure of the human conceptual system. *Cognitive Science*, 4(2):195–208, 1980.
- [38] Z. Liu and J. Stasko. Mental models, visual reasoning and interaction in information visualization: A top-down perspective. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):999–1008, 2010.
- [39] A. Ljubetic, F. Lapenta, H. Gradisar, I. Drobnak, J. Aupic, Z. Strmsek, D. Lainscek, I. Hafner-Bratkovic, A. Majerle, N. Krivec, M. Bencina, T. Pisanski, T. C. Velickovic, A. Round, J. M. Carazo, R. Melero, and R. Jerala. Design of coiled-coil protein-origami cages that self-assemble in vitro and in vivo. *Nature Biotechnology*, 35:1094 EP –, 2017.
- [40] E. D. Llano, H. Miao, M. E. G. Tobias Isenberg, I. Viola, and I. Barišić. A preview to Adenita: A software toolkit for the visualization and modeling of dna nanostructures. In *3rd Functional DNA Nanotechnology Workshop in Rome, Italy*, 2018.
- [41] A. Lu, C. J. Morris, J. Taylor, D. S. Ebert, C. Hansen, P. Rheingans, and M. Hartner. Illustrative interactive stipple rendering. *IEEE Transactions on Visualization and Computer Graphics*, 9(2):127–138, 2003.
- [42] D. Martín, G. Arroyo, A. Rodríguez, and T. Isenberg. A survey of digital stippling. *Computers & Graphics*, 67:24 – 44, 2017.
- [43] C. Menendez, J. Ordi, M. R. Ismail, P. J. Ventura, J. J. Aponte, E. Kahigwa, F. Font, and P. L. Alonso. The impact of placental malaria on gestational age and birth weight. *Journal of Infectious Diseases*, 181(5):1740–1745, 2000.
- [44] H. Miao, G. Mistelbauer, A. Karimov, A. Alansary, A. Davidson, D. F. Lloyd, M. Damodaram, L. Story, J. Hutter, J. V. Hajnal, M. Rutherford, B. Preim, B. Kainz, and M. E. Gröller. Placenta maps: In utero placental health assessment of the human fetus. *IEEE Transactions on Visualization and Computer Graphics*, 23(6):1612–1623, 2017.
- [45] H. Mohammed, A. K. Al-Awami, J. Beyer, C. Cali, P. Magistretti, H. Pfister, and M. Hadwiger. Abstractocyte: A visual tool for exploring nanoscale astroglial cells. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):853–861, 2018.
- [46] T. Munzner. A nested model for visualization design and validation. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):921–928, 2009.
- [47] NANO-D, Inria. SAMSON – Software for adaptive modeling and simulation of nanosystems. Website: <https://samson-connect.net>, 2019. Visited May 2019.

- [48] K. O'Donnell, T. O'Connor, and V. G. Glover. Prenatal stress and neurodevelopment of the child: Focus on the HPA axis and role of the placenta. *Developmental Neuroscience*, 31:285–292, 2009.
- [49] F. Praetorius and H. Dietz. Self-assembly of genetically encoded DNA-protein hybrid nanoscale shapes. *Science*, 355(6331), 2017.
- [50] E. Praun, H. Hoppe, M. Webb, and A. Finkelstein. Real-time hatching. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '01, page 581–586, New York, NY, USA, 2001. ACM.
- [51] P. Rautek, S. Bruckner, E. Gröller, and I. Viola. Illustrative visualization: New technology or useless tautology? *ACM SIGGRAPH Computer Graphics*, 42:4, 2008.
- [52] J. S. Risch. On the role of metaphor in information visualization. *Computing Research Repository*, abs/0809.0884, 2008.
- [53] J. C. Roberts. State of the art: Coordinated multiple views in exploratory visualization. In *Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV 2007)*, pages 61–71, 2007.
- [54] G. Robertson, S. K. Card, and J. D. Mackinlay. The cognitive coprocessor architecture for interactive user interfaces. In *Proceedings of the 2nd Annual ACM SIGGRAPH Symposium on User Interface Software and Technology*, UIST '89, pages 10–18, New York, NY, USA, 1989. ACM.
- [55] Roche. Biochemical Pathways. Website: <http://biochemical-pathways.com/#/map/1>, 2019. Visited in August 2019.
- [56] M. Rosenman and J. Gero. Modelling multiple views of design objects in a collaborative CAD environment. *Computer-Aided Design*, 28(3):193 – 205, 1996. Artificial Intelligence in Computer-Aided Design.
- [57] P. W. Rothemund. Folding DNA to create nanoscale shapes and patterns. *Nature*, 440(7082):297–302, 2006.
- [58] L. Saroul. *Surface Extraction and Flattening for Anatomical Visualization*. PhD thesis, Ecole Polytechnique Federale De Lausanne, 2006.
- [59] Schrödinger, LLC. Maestro 11: The completely reimaged all-purpose molecular modeling environment. Website: <https://www.schrodinger.com/maestro/>, 2019. Visited in August 2019.
- [60] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440, 2012.

- [61] J. Sweller. Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2):257–285, 1988.
- [62] J. G. Trafton, S. B. Trickett, and F. E. Mintz. Connecting internal and external representations: Spatial transformations of scientific visualizations. *Foundations of Science*, 10(1):89–106, 2005.
- [63] B. Tversky, J. B. Morrison, and M. Betrancourt. Animation: Can it facilitate? *International Journal of Human-Computer Studies*, 57(4):247–262, 2002.
- [64] M. van der Zwan, W. Lueks, H. Bekker, and T. Isenberg. Illustrative molecular visualization with continuous abstraction. *Computer Graphics Forum*, 30(3):683–690, 2011.
- [65] R. Veneziano, S. Ratanalert, K. Zhang, F. Zhang, H. Yan, W. Chiu, and M. Bathe. Designer nanoscale DNA assemblies programmed from the top down. *Science*, 352(6293):1534:1–1534:15, 2016.
- [66] I. Viola and M. E. Gröller. Smart visibility in visualization. In *Proceedings of the First Eurographics Conference on Computational Aesthetics in Graphics, Visualization and Imaging*, pages 209–216. The Eurographics Association, 2005.
- [67] I. Viola and T. Isenberg. Pondering the concept of abstraction in (illustrative) visualization. *IEEE Transactions on Visualization and Computer Graphics*, 24(9):2573–2588, 2018.
- [68] J. Watson and F. H. C. Crick. Molecular structure of nucleic acids; a structure for deoxyribose nucleic acid. *Nature*, 171:737–738, 1953.
- [69] A. Wolff. Drawing subway maps: A survey. *Informatik - Forschung und Entwicklung*, 22:23–44, 2007.
- [70] J. S. Yi, Y. A. Kang, and J. Stasko. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1224–1231, 2007.
- [71] J. Zander, T. Isenberg, S. Schlechtweg, and T. Strothotte. High quality hatching. *Computer Graphics Forum*, 23(3):421–430, 2004.

Publications

Placenta Maps: In Utero Placental Health Assessment of the Human Fetus

Haichao Miao, Gabriel Mistelbauer, Alexey Karimov, Amir Alansary, Alice Davidson, David F. A. Lloyd, Mellisa Damodaram, Lisa Story, Jana Hutter, Joseph V. Hajnal, Mary Rutherford, Bernhard Preim, Bernhard Kainz, and M. Eduard Gröller,

Abstract—The human placenta is essential for the supply of the fetus. To monitor the fetal development, imaging data is acquired using ultrasound (US). Although it is currently the gold-standard in fetal imaging, it might not capture certain abnormalities of the placenta. Magnetic resonance imaging (MRI) is a safe alternative for the in utero examination while acquiring the fetus data in higher detail. Nevertheless, there is currently no established procedure for assessing the condition of the placenta and consequently the fetal health. Due to maternal respiration and inherent movements of the fetus during examination, a quantitative assessment of the placenta requires fetal motion compensation, precise placenta segmentation and a standardized visualization, which are challenging tasks. Utilizing advanced motion compensation and automatic segmentation methods to extract the highly versatile shape of the placenta, we introduce a novel visualization technique that presents the fetal and maternal side of the placenta in a standardized way. Our approach enables physicians to explore the placenta even in utero. This establishes the basis for a comparative assessment of multiple placentas to analyze possible pathologic arrangements and to support the research and understanding of this vital organ. Additionally, we propose a three-dimensional structure-aware surface slicing technique in order to explore relevant regions inside the placenta. Finally, to survey the applicability of our approach, we consulted clinical experts in prenatal diagnostics and imaging. We received mainly positive feedback, especially the applicability of our technique for research purposes was appreciated.

Index Terms—Placenta, Fetal, Flattening, Structure-Aware Slicing, Peeling

1 INTRODUCTION

THE development and functions of the human placenta affect the fetus and, consequently, influence indicators of fetal health such as birth weight and growth [1], prematurity [2], and neuro-development [3]. There is also evidence that assessment of the placenta facilitates predictions of future health problems in adulthood [4]. This vital but barely studied organ controls the transmission of nutrients and hormones from the maternal to the fetal circulatory system and is essential for the fetal immune defense.

Placentas are currently primarily examined *ex vivo* by placing them on a flat table and assessing their structure [5]. The lack of standardized in utero representations of the placenta's structures renders clinical examinations and large population studies a challenging task. Figure 1 shows an *ex vivo* placenta augmented with a virtual compass for navigation purposes. The *ex vivo* placenta has predominately a disk-like, flat and round shape [5]. However, due to the limited space in the mother's womb and fetal movements, the placenta is deformed in utero, resulting in a highly diverse shape compared to the *ex vivo* placenta. As it grows

differently across subjects during gestation, its diversity is increased leading to an unpredictable random shape. This shape heterogeneity is a limiting factor when large population studies are carried out.

In order to enable *in utero* assessment of the human placenta we propose a novel visualization and deliver a proof of concept. On the one hand, conventional slice-by-slice volume inspection is a cumbersome task that requires mental matching of slices. On the other hand, volume rendering provides a good overview, but occlusion requires substantial interaction to inspect the placenta from different viewing angles. A standardized 2D representation would not only facilitate a faster inspection but also a comparison of multiple placentas across subjects. Motivated by standardized visualizations in medicine, such as the *Bull's eye plot* of coronary arteries [6] or CoWRadar of the Circle of Willis [7], we propose a novel standardized representation of the human placenta in utero, referred to as *placenta maps*, mimicking the familiar shape of the placenta *ex vivo*. The contributions of this paper are the following:

- H. Miao, A. Karimov and M. E. Gröller are with the Institute of Computer Graphics and Algorithms, TU Wien, Vienna, Austria.
- M. E. Gröller is also with the VRVis Research Center, Vienna, Austria.
- G. Mistelbauer and B. Preim are with the Department of Simulation and Graphics, Otto-von-Guericke University, Magdeburg, Germany.
- A. Alansary and B. Kainz are with the Department of Computing at the Imperial College, London, United Kingdom.
- A. Davidson, D. F. A. Lloyd, M. Damodaram, L. Story, J. Hutter, J. V. Hajnal, M. Rutherford, B. Kainz are with King's College London, United Kingdom.
- E-mail: miao@cg.tuwien.ac.at

Manuscript received April 19, 2005; revised August 26, 2015.

- An automatic identification of the fetal and maternal side of the placenta in utero,
- a structure-aware slicing approach to inspect the interior and exterior of the placenta in a well-defined fashion, and
- a standardized visual representation of the placenta showing the fetal and maternal side in two separate images allowing clinical experts a concise, comparative visual assessment.

In the context of large population studies of the human




Fig. 1. Two photographs of a healthy placenta extracted during a cesarean section. The fetal side is shown in (a) and the maternal side in (b). Since the placenta exhibits a disk-like shape when inspected on a table *ex vivo*, we decided for a circular visualization using a compass analogy for navigation, as overlaid.

placenta, the proposed approach could serve as a baseline visualization fostering in utero placenta research, as confirmed by our clinical collaborators.

2 RELATED WORK

In order to develop a standardized visualization of a three-dimensional object such as the placenta, a two-dimensional representation of the fetal as well as maternal side is a viable option. The placenta is inspected *ex vivo* on a table without any further interaction due to the removal of perspective occlusion. To achieve this, a projection from three to two dimensions is required and should be chosen carefully to preserve certain properties of the original object.

The shape of internal organs including the placenta can be highly irregular. Therefore, we need to globally analyze the input shape in terms of geometry. Saroul [8] investigates geometrical properties of two-dimensional surfaces in three-dimensional space. In particular, projection methods between two surfaces are inspected with respect to introduced distortions, e.g., lengths between points in the original surface may change in the projected surface. For medical experts the amount of such distortions is critical, as sizes and positions of anatomical features may not be measured with sufficient precision. Equiareal projections preserve areas of anatomical features across the entire placenta surface and conformal projections preserve orientations of the features with respect to angles. Only developable surfaces can be projected while preserving both areas and angles with isometric projections. Since the placenta surface is not developable in our case, we attempt to reduce both types of distortion, as they influence the interpretation of the data.

In the standard rendering pipeline, several voxels might contribute to a single image pixel of the final two-dimensional representation. To avoid this visual overload, clinical experts rely on two-dimensional slice views of the data. Complex structures of the human body, e.g., heart ventricles, vessel trees, bones in the rib cage region, may be incomprehensible on planar two-dimensional slices. Projection and reformation, as described by Preim and Botha [9], go a step further in terms of complexity and effectiveness. The authors identified three types of projections: anatomical unfolding, anatomical planar reformation or projection, and map projection.

Reformatted images can be computed to alleviate mental efforts and support examination [10], [11]. The reformation process relies on the principle that an arbitrary surface \mathcal{R} can be parametrized with two parameters u, v and then mapped to a rectilinear coordinate system, suitable for common display hardware. The reformation surface \mathcal{R} is often selected in such a way that it passes through relevant anatomical features. Interpretation of the reformation images requires the viewer to understand the location and the shape of \mathcal{R} . Therefore, standardization of the reformatted images is important in the medical domain. The technique described by Mistelbauer et al. [11] utilizes a so-called *anatomical layout* in order to connect various branches of a vessel tree. As a result, a medical expert can quickly identify the vessel currently being investigated. Moreover, one can seamlessly browse through the connected vessels.

Kretschmer et al. [10] propose anatomy-driven reformation (ADR), an approach that considers the overall shape of the assessed object for reformation. In this case, the ADR surface coincides with \mathcal{R} . They additionally compute offset surfaces in the positive and negative half-spaces of \mathcal{R} . These two surfaces represent boundaries of the sub-volume that is being reformatted. With the parametrization of \mathcal{R} , this technique tends to preserve the overall shape of the relevant object. The authors use the harmonic triangulation by Floater [12] and iteratively reduce distortions during an as-rigid-as-possible transformation. The step is required as this technique targets objects with different shapes, such as various bones in the human body. The initial harmonic projection reduces distortions in terms of the Dirichlet Energy. This can be thought of as getting a minimally stretched membrane over several protruded objects. While positions of the offset surface are locally optimized in order to minimize distortions, ADR does not explicitly minimize distortions between neighboring reformatted surfaces. This could be disadvantageous when investigating organs where connections between various tissues have to be preserved.

A standardized reformation of the myocardium is presented by Termeer et al. [6]. Representing the volume data with a volumetric *Bull's eye plot*, the technique allows the user to investigate the distribution of scar tissues within the myocardium. The myocardium is unfolded into a cylinder in such a way that all samples from a particular magnetic resonance imaging (MRI) slice, exhibiting a certain distance to the epicardium and endocardium, are arranged into a circle. While the structure of the myocardium may locally vary across different patients, the overall shape of it remains rather constant, leading to a standardized representation for the myocardium.

The visualization techniques described by Kanitsar et al. [13] are designed for reformatting the vascular structures in the human body. While taking the curvature of blood vessels into account, these methods differ in preserved properties. For example, the straightened curved planar reformation (CPR) preserves the length measured along the vascular structures and the stretched CPR leaves the curvatures unchanged. Since these methods require a centerline of the reformatted object they can only be applied to tubular-shaped objects. Vilanova et al. [14] discuss the unfolding of the colon surface. In order to achieve this, the centerline of the colon is extracted. Then, the surface of the

colon is parametrized using 2D coordinates and the unfolded image is generated using nonlinear ray casting. However, this method requires a well-defined centerline.

Krone et al. [15] apply well-known map projections from cartography to the domain of molecular visualization. Using gradient vector flow, they deform an arbitrary molecular surface to a sphere in order to facilitate map projections. Most of the techniques described in this section simplify the display of 3D anatomy in the 2D image space through reformation and projection. Since they cannot directly be applied to the human placenta due to its highly diverse shape, we introduce a flattened view of the placenta surface and allow the user to inspect its interior as well as exterior. Such advanced functionality provides crucial support for placenta research, as stated by our collaborating clinical experts.

3 MEDICAL BACKGROUND

Research on fetal MRI techniques [16] made it possible to image the placenta in greater detail than provided by routine ultrasound (US) examinations at week 12 and 20 of gestation. Besides US, MRI is the second non-invasive option to acquire image data from fetuses in utero, as other possibilities involve ionizing radiation or intervention. T2-weighted MRI has shown to exhibit good tissue contrast for fetal organs in several early studies [16], [17], but additional measurements like diffusion-weighted imaging and the identification of the maternal attachment area are becoming more popular [18].

MRI provides improved contrast and reduced noise in the acquired data of the fetal anatomy to detect abnormalities during pregnancy such as placental pathologies [19]. MRI is also considered to be safe after the first trimester [17] for 1.5 T [20] and 3 T [21] without the use of contrast agents, which may cause teratogenic effects on the fetus. It additionally enables researchers and clinical experts to analyze correlations between the childhood development and different prenatal abnormalities.

Motion corruption limits visual inspection and current visual exploration tools, as depicted in Figure 2. Advanced in utero studies of the placenta in 3D on large populations have not been conducted yet due to this problem [5]. In combination with fast sequences, such an imaging technique plays an essential role in the fetal diagnosis [16], in particular where US fails to provide sufficient data to diagnose certain conditions. MRI acquires high resolution slices from the uterus with a larger field-of-view in comparison to US. Furthermore, the only functional imaging option of US is Doppler imaging. MRI is of special interest because it can image the placenta as a whole, is able to acquire various functional data, and can show structure that is physically impossible to capture with US. However, there are still inter-slice artifacts, which consequently limit reliable diagnostics to individual slices. The observed motion of the placenta is of unpredictable nature, since it consists of maternal respiration movements, fetal movements and bowel movements. Predictable motion, like respiration, can be corrected with image navigator techniques or special MRI sequences [22]. But it is only possible to account for bowel and erratic fetal movements after image acquisition

using heavy oversampling of the input space, slice-to-volume registration, and super-resolution techniques [23], [24].

A quantitative assessment of the placenta is commonly done with respect to its volume. For example, Stevenson et al. [25] present a semi-automatic approach for measuring the volume of the placenta from motion-free 3D US data using a random walker algorithm. This method provides good inter-observer reproducibility, but requires some user interaction and several minutes per segmentation. Wang et al. [26] present an interactive method for placenta segmentation using MRI data, which requires an initial seed-point inside the placenta. Their approach performs well on a small cohort of six subjects, but shows a user-dependent variability in segmentation accuracy. In our work, we use the method presented by Alansary et al. [27] to achieve a motion-compensated segmentation of the placenta. This provides a solid baseline for our proposed standardized visual assessment approach. The achievable reconstruction quality using this approach has been evaluated in detail by Alansary et al. [28].

Image-based diagnostics are usually based on comparing to a memorized or actual *normative* anatomical appearance and on reproducible measurements on these images. Creating a standardized image is difficult for an organ like the placenta and requires specific methods discussed in this paper. The key benefit of *placenta maps* is that it is the first approach that allows the user comparative studies over large populations of placentas. Providing a method that is able to map the placenta into a normative space, primarily facilitates a comprehensive diagnostics of the placenta. Selected example pathologies, where a technique like *placenta maps* will have a direct impact, are twin-to-twin transfusion syndrome (TTTS), preeclampsia, preterm birth, and stillbirth. The visualization of relevant biomarkers, such as the overall appearance, texture, lesions, functional properties, and measurement of the attachment area, is eased through our approach. In clinical literature [29], [30], such biomarkers are considered to be vital measurements to assess the likelihood of complications during pregnancy.

Our collaborating clinical experts apply MRI to investigate TTTS. This syndrome is specific to monochorionic twins gestations, since the blood supply could be imbalanced and one of the twins is insufficiently supplied from the single shared placenta. In the case of TTTS the vessels branching off the two umbilical cords are interconnected with each other. TTTS is treated with endoscopic laser surgery to interrupt the vessel communications. Without intervention this syndrome is lethal to either one or both twins.

Cotyledons are small lobules that cover the entire maternal side, as shown in Figure 1b ex vivo. Ten to 40 cotyledons [5] are attached to the uterine wall and are responsible for the exchange of oxygen and nutrients between fetal and maternal blood. Because of the crucial function of cotyledons, they are of particular interest when examining the in utero placenta.

4 METHODOLOGY

In order to mimic the ex vivo assessment of the placenta we propose a novel visualization approach, outlined in Figure 3. It consists of the following four major automatic




Fig. 2. Three orthogonal cutting planes through a motion-corrupted dataset of T2-weighted MRI slices covering the uterus at a gestational age of 33 weeks are presented in (a). The motion-compensated reconstruction of the dataset [23] is shown in (b). The placenta is depicted by the white arrows.

steps: segmentation, slicing, splitting, and visual mapping. Firstly, the placenta is segmented and secondly, a distance field and its iso-surfaces are created. In the third step, the fetal and maternal side are automatically identified and fourthly, both three-dimensional meshes are projected into two-dimensional space. This workflow offers a standardized layout, and also favors inter-subject comparison. Subsequently, each step is explained in detail.

4.1 Segmentation

The input datasets were acquired by a 1.5 T Philips Achieva MRI system using a 32 channel cardiac array for signal reception. A total of four placentas with gestational age ranging from 24 to 37 weeks were motion-compensated by the technique of Kainz et al. [23] and segmented using the approach of Alansary et al. [27]. The resulting motion-compensated placenta masks are defined on a $0.75 \times 0.75 \times 0.75$ mm volumetric lattice and are shown in Figure 3 in the segmentation step.

We post-process the segmentation masks by applying a morphological opening with a spherical structuring element of radius five voxels to remove small local noise and objects. Afterwards, we perform morphological closing to fill small holes, again with a spherical structuring element but of radius ten voxels. The final mask of the placenta is then used in the subsequent steps.

4.2 Structure-Aware Slicing

Motivated by traditional axial, sagittal, and coronal slicing through the data for exploratory purposes, we apply this concept to the structure of the placenta itself. Instead of slicing along a defined axis, we compute the Euclidean distance field from the surface \mathcal{R} of the placenta using the method described by Danielsson [31].

We then compute offset surfaces of \mathcal{R} at discrete distances of 1 mm. Since these surfaces define layers, we denote \mathcal{R} with \mathcal{L}_0 , the surfaces inside the placenta with positive indices




Fig. 3. Our automatic pipeline consists of four steps. *Segmentation* generates the mask for the placenta. *Slicing* creates the layers based on the placenta structure. For demonstration purposes we slice every 2 mm. *Splitting* separates the layers into a fetal and maternal side and each is then flattened. Finally, *Visual Mapping* creates a standardized view of the isolated placenta. Each location on the fetal and maternal side can be addressed with latitude (1 – 3) and longitude ($1^\circ - 360^\circ$) coordinates using the compass overlay.

and the ones outside with negative indices. The layers are illustrated in Figure 4. Slicing into the placenta is practical, as it reveals the interior layer by layer. However, we also provide means to display the context of the placenta by slicing outside of the placenta. With outside slicing, we can visualize the uterine wall on the maternal side and the amniotic cavity on the fetal side with respect to the placenta structure. The slicing leads to a set of layers \mathcal{L}_i , each consisting of a set V of vertices $v_k \in \mathbb{R}^3$ and faces $f_j \in F$. The image of the slicing step in Figure 3 shows the distance field on the left as well as the internal layers on the right.




Fig. 4. The layers of the placenta. First, the layers $\mathcal{L}_{-9} \dots \mathcal{L}_5$ are split into a blue fetal side ($\mathcal{F}_{-9} \dots \mathcal{F}_5$) and a yellow maternal side ($\mathcal{M}_{-9} \dots \mathcal{M}_5$). Layers are shown in 3D and at the bottom in the standardized, flattened view. In 3D, it can be seen that the shape gets smaller and smaller as layers are removed. \mathcal{F}_{-9} shows the two twins. The structural appearance in \mathcal{M}_0 (maternal) indicates the cotyledons.

In contrast to our approach, the ADR technique by Kretschmer et al. [10] uses the normal vectors of a surface to compute its offset surfaces. However, this possibly leads to self-intersections and not well-defined surfaces, especially in regions with high curvature or in case of large offsets. Using a distance field, we avoid these types of problems.

Since we flatten each layer in the subsequent step and provide a two-dimensional map, the user can scroll through the different layers of the placenta in the usual fashion. This offers the possibility to see structures at a certain depth or height, such as vessels, cotyledons or diseased tissues.

4.3 Splitting

We divide each extracted layer into two parts that reflect the fetal and maternal side as shown in the splitting step in Figure 3. We refer to this fetal and maternal side detection as simply splitting.

As we rely on a coarse orientation of the placental shape, we perform principal component analysis (PCA). The basic idea is to consider the vertices of the layer \mathcal{L}_i as a set of 3D points that is projected into a linear subspace. This resembles fitting a plane to the surface, where the two axes of the plane describe the direction of the highest and the second highest variance in terms of mesh point positions. The best fitting plane is found by using a linear least-square fitting method, which minimizes the sum of squared distances from all points to the plane, as shown in the image of the splitting step in Figure 3. As the normal vector \mathbf{p} of this plane represents the direction of the least variance, it reflects the thickness of the placenta. Vector \mathbf{q} is defined as the opposite of \mathbf{p} .

In order to split \mathcal{L}_i into two parts \mathcal{F}_i and \mathcal{M}_i , we first find two faces f_A and f_B such that $f_A \in \mathcal{F}_i$ and $f_B \in \mathcal{M}_i$. They are found by intersecting the mesh with a line defined

by the vectors \mathbf{p} and \mathbf{q} and a point O at the or close to the center of the placenta. We set the origin of our coordinate system to this point. Consequently, O has to be located inside the placenta at a central position. We define O as a single member of the skeleton S that minimizes the sum of distances to the remaining skeleton points. We apply the voxel-based skeletonization approach by Lee et al. [32] that uses topological thinning to extract the skeleton S of an object. From the skeleton voxels we find O as follows:

$$O = \arg \min_{X \in S} \sum_{R \in S} \|X - R\|_2. \quad (1)$$

This formula is based on the geometric median, however, it minimizes the sum of distances to a member of the skeleton. As O is located on the skeleton of the placenta it is ensured that O is always located inside the placenta. This is not guaranteed if the centroid or the center of mass is used. They could lie outside the shape if the in utero placenta is concave, which is rather common.

To separate \mathcal{L}_i into the two parts \mathcal{F}_i and \mathcal{M}_i corresponding to the fetal and maternal side, we combine the *silhouetteness* property of faces with their breadth-first traversal. We construct a graph where vertices are mesh faces and edges connect adjacent faces. The faces f_A and f_B are the starting points of the breadth-first traversal. During the traversal, all neighboring faces of f_j are visited first and with each face, the cost of 1 is added to include the approximate distance to f_A and f_B . This way we guarantee that each part is a connected set of faces, which is a requirement for the flattening. For each face f_j we compute two costs, i.e., one for belonging to the fetal part $c_{\mathcal{F}}(f_j)$ and one for belonging

to the maternal part $c_{\mathcal{M}}(f_j)$:

$$f_j \in \begin{cases} \mathcal{F}, & \text{if } c_{\mathcal{F}}(f_j) < c_{\mathcal{M}}(f_j) \\ \mathcal{M}, & \text{otherwise} \end{cases}, \quad (2)$$

$$c_{\mathcal{F}}(f_j) = c_{\mathcal{F}}(f_{j-1}) + s(\mathbf{n}(f_j), \mathbf{q}) + 1, \quad (3)$$

$$c_{\mathcal{M}}(f_j) = c_{\mathcal{M}}(f_{j-1}) + s(\mathbf{n}(f_j), \mathbf{p}) + 1, \quad (4)$$

$$c_{\mathcal{F}}(f_0 = f_A) = 0, \quad c_{\mathcal{M}}(f_0 = f_B) = 0, \quad (5)$$

$$s(\mathbf{n}(f_j), \mathbf{v}) = \angle(\mathbf{n}(f_j), \mathbf{v}). \quad (6)$$

where $\mathbf{n}(f_j)$ denotes the normal to the face f_j and $c_{\mathcal{F}}(f_{j-1})$ and $c_{\mathcal{M}}(f_{j-1})$ are the costs of the previously traversed face f_{j-1} . The underlying idea of the silhouetteness term $s(\cdot)$ is that f_j is more likely to belong to \mathcal{F} if $\mathbf{n}(f_j)$ is aligned with \mathbf{q} . Therefore, we use the angle between $\mathbf{n}(f_j)$ and \mathbf{q} , which is motivated by the common approach to detect silhouettes in rendering. The silhouetteness term $s(\cdot)$ measures the angle in radians. The cost is increased by 1 if $\mathbf{n}(f_j)$ points in the same direction as \mathbf{q} and by $1 + \pi$ if the two vectors have opposite directions. If only the *silhouetteness* term is considered, local curvature changes may cause the assignment of the faces to the wrong side due to the change of face normals. By accumulating the costs, we ensure that f_j is connected to f_0 inside either \mathcal{F}_i or \mathcal{M}_i . This means that the more distant f_j is to f_0 , the higher is the cost. These two terms ensure that crossing from one side to another is very costly. This splitting of the initial surface into two parts is shown in the splitting step of [Figure 3](#) and in [Figure 4](#).

We compute which of the sides is the fetal and which is the maternal one. Anatomically, the fetal side is proximal, i.e., facing towards the body center, and the maternal side is distal, i.e., facing away from the body center. The proximity is estimated with respect to the volume center. We can safely assume that the fetus is located in the volume center, as it coincides with the region of interest during the data acquisition of the fetal MRI.

In the case that the principal component analysis (PCA) fails, e.g. if the input shape is sphere-like, which is unlikely, then we suggest a fallback solution: the user can manually specify f_A and f_B . The remaining splitting step would still be automatic.

Flattening

The surface of the placenta can exhibit large local curvature. This hampers its investigation in conventional slice views commonly used by the clinical experts. Following the standard procedure of physical examination of the placenta *ex vivo*, we ease the exploration of the placenta in utero by flattening it. As a result, we create a standardized view of the placenta, as shown at the bottom of [Figure 4](#).

Our method is based on the fact that the *ex vivo* placenta can be physically deformed into a disk shape (see [Figure 1](#)). The clinical experts distinguish between the fetal and the maternal side of the placenta. Therefore, we split the original placenta surface into these two parts and flatten each part separately. Because of the disk shape, we do not need additional transformations after we flatten the surface with the *mean value coordinates* approach of Floater [33]. In this flattening, the original in utero shape information is lost. However, the specific placenta shape is functionally not important [5]. The surface \mathcal{P} must be open, which we achieve

by splitting the entire layer into two parts as described in [subsection 4.3](#). The boundary points of the surface \mathcal{P} are mapped to a disc. Projections of the remaining points are then determined inside the disk via the convex combination mapping. The convex combination mapping utilizes the mean value coordinates of each point in the convex basis formed by neighboring points [33]. After the flattening we refer to the flattened fetal and maternal side as \mathcal{F}_i' and \mathcal{M}_i' respectively. The faces that form the corresponding meshes are denoted as f_j' .

In [Figure 5](#) we map a checkerboard pattern onto the placenta surface to show the angular distortion between the original and the flattened meshes. To quantify the area distortions, we compute the relative face-area deviation and color-code it. The area deviation of a face f_j is thereby computed as follows:

$$\eta(f_j') = \frac{|A(f_j') - A(f_j)|}{A(f_j)}, \quad (7)$$

where $A(f_j)$ denotes the area of f_j .

4.4 Visual Mapping

A placenta map is a visualization that consists of two views, as shown in [Figure 4](#). The first view is a 3D visualization of placenta layers and the second one is the separate flattened reformatted visualization of the fetal and maternal sides. In the last image of [Figure 3](#), we overlay a compass onto the maternal side for navigation purposes. The compass divides each side into four quadrants. Each quadrant is specified by a cardinal direction (north, east, south and west) for a rough guidance. For a more accurate addressing, latitude ($1 - 3$) and longitude (between 1° and 360°) coordinates are used. For example, in the first two images of [Figure 5](#) the red part is marked at $(2, 315^\circ - 45^\circ)$. This specifies a latitude of 2 with a longitude between 315° and 45° . The compass is thereby standardized and overlaid over the 3D and 2D view. This enables experts to mark diseased portions on the placenta utilizing a link between the 3D and 2D view.

This way we present the fetal and maternal side simultaneously without occlusion as shown in [Figure 5](#). By traversing all layers, the clinical expert can quickly browse through the placenta and investigate its interior with an easier navigation in reformatted and standardized images than in conventional slice views.

Using standardized flattened views, the comparison of multiple placentas becomes feasible. The overlaid compass can be reoriented manually to match several placentas. The comparison of placentas from different patients or multiple time-stamps of the same patient is an interesting extension to existing methods in placenta research. It may reveal new insights about diseased regions and their influence during the gestation. In order to enable the comparison, we juxtapose the flattened placenta images. Using the compass overlay, the expert can match regions on the flattened images from different subjects and time-stamps. With slice views or volume rendered images, such a comparison is rather difficult, because the placentas vary significantly in shape. Thus, the mapping between two or more placentas is too complicated to easily support a mental model.

Clinical experts can potentially investigate diseases, such as TTTS, and track vessels on the fetal side. Having just





Fig. 5. Flattening of the fetal and maternal side. The images on the left show the placenta maps of the fetal as well as maternal side. A compass is projected onto both sides in 3D and 2D. Two areas highlighted in red correspond to each other. In the central images, a checkerboard pattern is orthogonally projected along p to demonstrate the angular distortion between the 3D and 2D views. p is then aligned with the direction to the viewer. On the right side, the area distortion $\eta(\cdot)$ is shown per face, with colors  representing the intervals $[0; 0.2)$, $[0.2, 0.4)$, $[0.4, 0.6)$, $[0.6, 0.8)$, $[0.8, +\infty)$ respectively.

conventional slice views, the expert would need to build up a mental model of the placenta and its vascular structures. Our technique shows in 2D and at once the entire fetal side where the vascular structures are located. The vessels inside the placenta could be easily investigated with our structure-aware slicing method.

5 IMPLEMENTATION

Placenta maps use the Qt, the ITK, and the VTK framework, executed on an Intel Core i7 CPU with an NVidia GeForce 980 Ti GPU. The structure-aware slicing and splitting are performed on the CPU. For the mean value coordinates technique and the PCA we use the CGAL library [34]. In order to provide an interactive investigation of the data, the layers of the placenta are automatically flattened beforehand.

6 RESULTS AND DISCUSSION

In this section, we demonstrate our approach on three datasets, shown in Figure 4, Figure 6 and Figure 7. Each dataset is discussed and we explain how potential findings can be achieved using placenta maps. We also discuss how to simulate cotyledons and how our slicing displays them.

The placenta of a subject at a gestational age of 26 weeks having monochorionic twins and the diagnosis of TTTS is shown in Figure 4. The top image depicts the segmented placenta and an exploded view of our structure-aware slicing method with the split fetal and maternal side. Layer \mathcal{F}_{-9} displays both twins, the first twin is located at the north in the placenta map and the second twin is partially visible at the south. As the twins are visible on the fetal side, this suggests that our automatic splitting technique correctly classifies both, the fetal and maternal side. \mathcal{F}_0 and \mathcal{M}_0 show the initial surface of the segmented placenta. Within our proposed placenta map in \mathcal{F}_0 the suspected umbilical cord insertion is visible at $(3, 160^\circ)$. In the maternal layer \mathcal{M}_0 ,

the cotyledons are depicted, though in low detail due to yet insufficient quality of the data acquisition. \mathcal{F}_5 and \mathcal{M}_5 display internal tissues of the placenta. The placenta maps allow physicians to monitor the development of the in utero placenta in a standardized way.

Another subject at a gestational age of 30 weeks is presented in Figure 6. Earlier occurrence of the TTTS is suspected in this subject, but laser surgery was not applied. One twin died in week 23, but the other survived. The uterine wall is partially depicted in \mathcal{M}_{-10} . The surviving twin is visible on the entire placenta map in \mathcal{F}_{-10} . Amniotic fluid appears as bright matter, exhibiting high intensity values and surrounding the twins. In \mathcal{F}_{-10} , showing the amniotic cavity, the umbilical cord of the surviving twin can be seen in the East quadrant, precisely at $(2, 45^\circ - 160^\circ)$. The insertion of the cord on the fetal side at \mathcal{F}_0 is at $(2, 45^\circ)$.

The last subject is at a gestational age of 37 weeks and shown in Figure 7. Although the placenta is healthy, the subject suffered a cytomegalovirus infection, causing microcephaly in the fetus. The anatomical context at \mathcal{F}_{-10} displays the fetus. This indicates the correct automatic splitting of the placenta into the fetal and maternal side. At \mathcal{F}_0 the insertion of the umbilical cord to the fetal side is at $(2, 140^\circ - 150^\circ)$. Cotyledons appear on the maternal side \mathcal{M}_0 . \mathcal{F}_3 and \mathcal{M}_3 present slices inside the placenta. Here, vessels branching from the umbilical cord can be observed at $(2, 160^\circ)$ in \mathcal{F}_3 .

Figure 8 depicts artificial cotyledons. Since the fetus is focused during the data acquisition, the signal-to-noise ratio of the placenta is low. This makes the visual assessment of the placenta a challenging task, especially of structures with low contrast, such as the cotyledons. Our collaborators explicitly want cotyledons in the visualizations to gain an understanding of the uterine attachment when using placenta maps. In order to demonstrate such functionality, we artificially created cotyledons in a concept study. Based

on the ex vivo appearance of cotyledons, we modeled them as spheres with center points on the maternal side. The intensity values of the voxels inside a cotyledon are defined as $I = [1 - (d/r)^s] \cdot 2000$, where d is the Euclidean distance to the center of the sphere and s controls the fall-off in intensity values. We empirically set radius r to 25 voxels and s to 1.7. The pattern of the cotyledons remains stable with respect to the layer depth. Hence, our technique shows internal structures with rather small distortions between adjacent layers.

7 EVALUATION

To assess the applicability and acceptance of the proposed technique to in utero examination of the placenta, we conducted a survey. It was based on questions with answers on a five-point Likert-scale and open questions. The questions were grouped into the following three core aspects of our technique: structure-aware slicing, comparability, and flattening. The first part of the survey consists of pre-recorded (R) videos that demonstrate the technique to the clinical experts. The videos showed the three aspects separately and the experts rated their usefulness. To assess the acceptance of our technique, we designed the second part of the survey to be interactive (I). The experts tested our tool with multiple datasets and rated the usability of the same three aspects. Finally, the participants assessed the overall readability of placenta maps. We asked summary questions regarding the applicability of our technique to placenta research and also to clinical routine. Figure 9 depicts the results of the survey, where the answers from the pre-recorded (R) and the interactive (I) presentation of the aspects are separated. For the pre-recorded part, we received four responses from clinical experts, where three are radiologists specialized in prenatal diagnostics and one is a pediatrician. The interactive part of the survey was conducted by three experts.

Our proposed **structure-aware slicing** was rated by the experts with a score of 4.4 (R) and 4.0 (I). The experts especially appreciate the ability to display the isolated placenta and the slicing in order to see the internal tissues layer by layer with little mental effort. We assume, this offers the experts an efficient way to go through the different tissue layers. One expert stated that our slicing approach is very promising for tracking blood vessel, e.g., in case of TTTS.

The **comparability** of our standardized visualization got the best scores among the three aspects, with 4.3 (R) and 5.0 (I). The increase of the score after the experts tested the tool shows that this is the most valued aspect of placenta maps. Also, the qualitative feedback so far was clearly indicating that this feature is very important if a large number of placentas has to be investigated within a study. Ultimately, we aim to provide quantitative measurements and comparisons of the fetal and maternal sides.

We get a 2D view of the fetal and maternal side with **flattening** at the expense of accuracy, since each side is distorted as shown in Figure 5. This feature was rated with scores of 4.2 (R) and 4.3 (I). From the earlier qualitative feedback of one of the experts, we know that this feature can be very useful for placenta research, as it provides a fast overview of the fetal and maternal side. We quantify the introduced distortions later in Figure 10 to validate our




Fig. 6. Subject at a gestational age of 30 weeks with twins. The placenta map \mathcal{F}_{-10} shows the context of the placenta where the surviving twin is visible. The amniotic fluid has high intensity values, mapped to white. The layer \mathcal{F}_0 displays the surface of the fetal side, where one umbilical cord is visible as a dark spot. \mathcal{F}_3 and \mathcal{M}_3 show the internal tissues of the placenta when three layers are already removed. The three-dimensional views clearly show the decrease of the layers in size caused by their successive removal.




Fig. 7. Subject at a gestational age of 37 weeks with a healthy placenta, but suffering from an infection. \mathcal{F}_{-10} and \mathcal{M}_{-10} display the context of the placenta and clearly delineate the fetus in the amniotic cavity in \mathcal{F}_{-10} . This indicates that our splitting approach correctly detected the fetal and maternal side. \mathcal{M}_0 shows the surface of the placenta with cotyledons in low detail on the maternal side. \mathcal{F}_3 and \mathcal{M}_3 display the internal placenta tissues, after removing three layers.




Fig. 8. Artificial cotyledons modeled as spheres across the maternal side of the placenta. The different layers demonstrate the depth of the artificial cotyledons.

flattening technique. To satisfy the clinical requirements concerning accuracy, we additionally display axial slices that are linked to the placenta maps and the 3D view, as shown in the segmentation step of Figure 3.

The **overall readability** of our method received a score of 4.3, but could be improved. We suspect that the readability is attributed to our compass overlay, which is a novel way to address the heterogeneous placenta surface and, as stated by one of the experts, needs some familiarization time.

Regarding the **clinical applicability**, our technique got a score of 4.5. While the radiologists in general gave better ratings, the lowest rating of three was given by the pediatrician. We assume that this is attributable to the fact that he is less experienced with reading MRI data than radiologists.

The **research applicability** was rated with the best scores of five from all participants. According to one of our experts, the attachment to the uterine wall is of potential interest. However, a large number of placentas will have to be investigated, which would be feasible with placenta maps. The high rating gives a clear indication that our technique has potential to be applied in the research of placentas, which has been the main motivation of our approach.

In the open questions part of the survey, we asked the participants about their preferences, shortcomings of the technique, and further suggestions. We inquired which **different diseases and conditions** could be investigated with placenta maps. TTTS, placental infarction, previous laser treatment, and placenta praevia has been mentioned by two participants.

Regarding the **preferences**, two experts liked the possibility to compare multiple placentas enabled by our standardized, flattened views. Two participants appreciated the compass overlay on the placenta and consider this feature useful to visualize diseased portions.

Regarding the **shortcomings**, all experts agreed that the signal-to-noise ratio of the original data is insufficient for perceiving the features of interest, such as cotyledons and vessels. The primary cause of this issue is the used MRI sequence that does not target the placenta tissues specifically.

We also got **suggestions for improvement**. One expert recommends to enhance the segmentation result before mapping it to a 2D structure. This concerns the quality of




Fig. 9. The outcome of the survey. We differentiate between a pre-recorded (R) and an interactive (I) part. The scores are from one being poor to five being excellent.

the data, but not our technique. The clinical experts were convinced that with an improved signal-to-noise ratio of the data our technique would be highly relevant for research purposes and also for clinical use cases. The participants consider it useful to segment the umbilical cord and the cotyledons. The measurements of linear sizes were also an interesting suggestion, since they could improve the compass overlay to be applicable for quantitative measurements.

The flattening inevitably introduces some distortions. After interacting with our tool, we interviewed three experts regarding the acceptance of these distortions for clinical and exploratory usage. In the **clinical use case** the distortions are regarded as very acceptable, especially if the inspection only concerns the quantification of areas and counting cotyledons. The participants stated that our technique is much more useful than cutting planes. However, the simultaneous display of the undistorted view is necessary to provide the relationship to the unchanged anatomy. In the **exploratory use case**, the distortions are definitely acceptable for the intended purposes. One expert stated, that this technique is exactly what clinical research of the placenta currently needs. An appropriate training is deemed necessary in order to fully understand the placenta maps. In conclusion, it seems that the benefits of standardization, comparability, and fast assessment of unusual placentas outweighs the distortions introduced by our approach.

We quantified the distortions in all presented datasets using the distribution of **area deviations** as metric. For all faces $f_j' \in \mathcal{L}_0'$ we compute the area deviation $\eta(f_j')$ using Equation 7. We then aggregate the results into bins of size 0.2 and plot the resulting histogram. As shown in Figure 10, in all four datasets at least 43% of projected faces have areas that deviate less than 20%. This clearly indicates that our technique only introduces minor distortions.

8 LIMITATIONS

Our technique has some limitations that are discussed in the following. First of all, it slices the placenta in a discrete manner. A continuous layer definition is desirable in order to display the details between the layers. Such a continuous definition, however, was not in the focus of this work.

In general, our distance field-based solution extracts the offset surfaces while avoiding self-intersections. However, the topology of deeply located layers changes as the placenta thickness varies. By correcting the distance field using the information derived from the medial axis surface, it could be possible to preserve the topology. However, this requires further exploration and investigation.




Fig. 10. Distribution of area deviation $\eta(\cdot)$ of all presented datasets. Most faces exhibit a low distortion in the interval of $[0, 0.2)$.

In our splitting step, the distance term of the cost function is increased incrementally with every face in the breadth-first traversal. Currently, the cost function does not consider the area of a face, which is considered a future aspect. As confirmed by the positive feedback of our collaborating clinical experts, our approach delivered satisfactory results for all presented datasets.

The ultimate goal of placenta maps is to aid researchers in the examination of placentas and to compare them in large population studies. With placenta maps, this now seems feasible. It is still an open topic how the comparison should be performed. Our technique cannot yet be utilized to create precise matches between multiple placentas, as there are currently no landmarks detected to align the compass overlay. Techniques from comparative visualization can be used to improve the compass overlay functionality with respect to multiple placentas. To validate and further improve our technique, larger user studies have yet to be carried out.

9 CONCLUSION AND FUTURE WORK

In this work, we presented placenta maps, a novel technique to visualize the in utero placenta in MRI data. It features an automatic reformation of the placenta into the fetal and maternal side, which creates occlusion-free 2D views respectively. In addition, it enables structure-aware slicing, which shows the internal tissues, and also the anatomical context of the placenta. The positive outcome of our conducted survey indicates that our technique may open new possibilities for in utero placenta examinations.

We see several directions for future work. One of the most interesting extensions is to integrate secondary tissue information, for example obtained by diffusion weighted imaging. According to the experts, the ability to compare the entire surface of multiple placentas is highly relevant and potentially will increase our understanding of the placental attachment to the uterine wall.

ACKNOWLEDGMENTS

This project has been supported by the Vienna Science and Technology Fund (WWTF) through project VRG11-010, by EC Marie Curie Career Integration Grant through project PCIG13-GA-2013-618680, by Nvidia (Tesla K40 donation), by the EPSRC award EP/N024494/1, and by the National Institute for Health Research (NIHR) Biomedical Research Centre based at Guy’s and St Thomas’ NHS Foundation Trust and King’s College London. The views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR or the Department of Health.

REFERENCES

- [1] C. Menendez, J. Ordi, M. R. Ismail, P. J. Ventura, J. J. Aponte, E. Kahigwa, F. Font, and P. L. Alonso, "The impact of placental malaria on gestational age and birth weight," *Journal of Infectious Diseases*, vol. 181, no. 5, pp. 1740–1745, 2000.
- [2] F. Arias, L. Rodriguez, S. C. Rayne, and F. T. Kraus, "Maternal placental vasculopathy and infection: Two distinct subgroups among patients with preterm labor and preterm ruptured membranes," *American Journal of Obstetrics and Gynecology*, vol. 168, no. 2, pp. 585–591, 1993.
- [3] K. O'Donnell, T. O'Connor, and V. G. Glover, "Prenatal stress and neurodevelopment of the child: Focus on the HPA axis and role of the placenta," *Developmental Neuroscience*, vol. 31, pp. 285–292, 2009.
- [4] D. J. Barker, A. R. Bull, C. Osmond, and S. J. Simmonds, "Fetal and placental size and risk of hypertension in adult life," *British Medical Journal*, vol. 301, no. 6746, pp. 259–262, 1990.
- [5] K. Benirschke, G. J. Burton, and R. N. Baergen, *Pathology of the Human Placenta*. Springer Science & Business Media, 2012.
- [6] M. Termeer, J. O. Bescós, M. Breeuwer, A. Vilanova, F. Gerritsen, and E. Gröller, "CoViCAD: Comprehensive visualization of coronary artery disease," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1632–1639, 2007.
- [7] H. Miao, G. Mistelbauer, C. Nasel, and M. E. Gröller, "Visual quantification of the Circle of Willis: An automated identification and standardized representation," *Computer Graphics Forum*, 2016.
- [8] L. Saroul, "Surface extraction and flattening for anatomical visualization," PhD thesis, Ecole Polytechnique Federale De Lausanne, 2006.
- [9] B. Preim and C. P. Botha, *Visual Computing for Medicine*. Elsevier, 2013.
- [10] J. Kretschmer, G. Soza, C. Tietjen, M. Suehling, B. Preim, and M. Stamminger, "ADR - Anatomy-Driven Reformation," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 12, pp. 2496–2505, 2014.
- [11] G. Mistelbauer, A. Morar, A. Varchola, R. Scherthner, I. Baclija, A. Köchl, A. Kanitsar, S. Bruckner, and E. Gröller, "Vessel Visualization using curvicircular feature aggregation," *Computer Graphics Forum*, vol. 32, no. 3pt2, pp. 231–240, 2013.
- [12] M. S. Floater, "Parametrization and smooth approximation of surface triangulations," *Computer Aided Geometric Design*, vol. 14, no. 3, pp. 231–250, 1997.
- [13] A. Kanitsar, D. Fleischmann, R. Wegenkittl, P. Felkel, and M. E. Gröller, "CPR - Curved Planar Reformation," in *Proceedings of IEEE Visualization, 2002*, pp. 37–44.
- [14] A. Vilanova, R. Wegenkittl, A. König, and M. E. Gröller, "Nonlinear virtual colon unfolding," in *Proceedings of IEEE Visualization, 2001*, pp. 411–418.
- [15] M. Krone, F. Frieß, K. Scharnowski, G. Reina, S. Fademrecht, T. Kulschewski, J. Pleiss, and T. Ertl, "Molecular surface maps," *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 1, 2016.
- [16] B. Ertl-Wagner, A. Lienemann, A. Strauss, and M. F. Reiser, "Fetal magnetic resonance imaging: indications, technique, anatomical considerations and a review of fetal abnormalities," *European Radiology*, vol. 12, no. 8, pp. 1931–1940, 2002.
- [17] M. C. Frates, A. J. Kumar, C. B. Benson, V. L. Ward, and C. M. Tempny, "Fetal anomalies: Comparison of MR imaging and US for diagnosis," *Radiology*, vol. 232, no. 2, pp. 398–404, 2004.
- [18] L. Manganaro, F. Fierro, A. Tomei, L. La Barbera, S. Savelli, P. Sollazzo, M. Eleonora Sergi, V. Vinci, L. Balleio, and M. Marini, "MRI and DWI: feasibility of DWI and ADC maps in the evaluation of placental changes during gestation," *Prenatal Diagnosis*, vol. 30, no. 12-13, pp. 1178–1184, 2010.
- [19] N. Linduska, S. Dekan, A. Messerschmidt, G. Kasprian, P. Brugger, K. Chalubinski, M. Weber, and D. Prayer, "Placental pathologies in fetal MRI with pathohistological correlation," *Placenta*, vol. 30, no. 6, pp. 555–559, 2009.
- [20] M. Bouyssi-Kobar, A. J. du Plessis, R. L. Robertson, and C. Limperopoulos, "Fetal magnetic resonance imaging: exposure times and functional outcomes at preschool age," *Pediatric Radiology*, vol. 45, no. 12, pp. 1823–1830, 2015.
- [21] M. M. Cannie, F. De Keyzer, S. Van Laere, A. Leus, J. de Mey, C. Fourneau, F. De Ridder, T. Van Caueren, I. Willekens, and J. C. Jani, "Potential heating effect in the gravid uterus by using 3-T MR imaging protocols: Experimental study in miniature pigs," *Radiology*, vol. 279, no. 3, pp. 754–761, Dec 2016.
- [22] C. Malamateniou, S. J. Malik, S. J. Counsell, J. M. Allsop, A. K. McGuinness, T. Hayat, K. Broadhouse, R. G. Nunes, A. M. Ederies, J. V. Hajnal, and M. A. Rutherford, "Motion-compensation techniques in neonatal and fetal MR imaging," *American Journal of Neuroradiology*, vol. 34, no. 6, pp. 1124–1136, 2012.
- [23] B. Kainz, A. Alansary, C. Malamateniou, K. Keradren, M. Rutherford, J. V. Hajnal, and D. Rueckert, *Flexible Reconstruction and Correction of Unpredictable Motion from Stacks of 2D Images*. Springer International Publishing, 2015, pp. 555–562.
- [24] B. Kainz, M. Steinberger, W. Wein, M. Murgasova, C. Malamateniou, K. Keradren, T. Torsney-Weir, K., P. Aljabar, M. Rutherford, J. Hajnal, and D. Rueckert, "Fast volume reconstruction from motion corrupted stacks of 2D slices," *IEEE Transactions on Medical Imaging*, vol. 34, no. 9, pp. 1901–1913, 2015.
- [25] G. N. Stevenson, S. L. Collins, J. Ding, L. Impey, and J. A. Noble, "3-D ultrasound segmentation of the placenta using the random walker algorithm: Reliability and agreement," *Ultrasound in Medicine & Biology*, vol. 41, no. 12, pp. 3182–3193, 2015.
- [26] G. Wang, M. A. Zuluaga, R. Pratt, M. Aertsen, A. L. David, J. Deprest, T. Vercauteren, and S. Ourselin, "Slic-Seg: slice-by-slice segmentation propagation of the placenta in fetal MRI using one-plane scribbles and online learning," in *Proceedings of Medical Image Computing and Computer-Assisted Intervention, 2015*, pp. 29–37.
- [27] A. Alansary, K. Kamnitsas, M. Rajchl, A. Davidson, R. Khlebnikov, C. Malamateniou, M. Rutherford, J. Hajnal, B. Glocker, D. Rueckert, and B. Kainz, "Fast fully automatic segmentation of the human placenta from motion corrupted MRI," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2016, pp. 589–597.
- [28] A. Alansary, B. Kainz, M. Rajchl, M. Murgasova, M. Damodaram, D. F. Lloyd, A. Davidson, S. G. McDonagh, M. Rutherford, J. V. Hajnal, and D. Rueckert, "PVR: Patch-to-Volume Reconstruction for large area motion correction of fetal MRI," 2016, preprint available on arXiv:1611.07289.
- [29] R. Redline, D. Heller, S. Keating, and J. Kingdom, "Placental diagnostic criteria and clinical correlation - a workshop report," *Placenta*, vol. 26, Supplement, pp. S114 – S117, 2005.
- [30] H. S. Wong, Y. K. Cheung, J. Zuccollo, J. Tait, and K. C. Pringle, "Evaluation of sonographic diagnostic criteria for placenta accreta," *Journal of Clinical Ultrasound*, vol. 36, no. 9, pp. 551–559, 2008.
- [31] P.-E. Danielsson, "Euclidean distance mapping," *Computer Graphics and Image Processing*, vol. 14, no. 3, pp. 227–248, 1980.
- [32] T.-C. Lee, R. L. Kashyap, and C.-N. Chu, "Building skeleton models via 3-D medial surface/axis thinning algorithms," *Graphical Models and Image Processing*, vol. 56, no. 6, pp. 462–478, 1994.
- [33] M. S. Floater, "Mean value coordinates," *Computer Aided Geometric Design*, vol. 20, no. 1, pp. 19–27, 2003.
- [34] The CGAL Project, *CGAL User and Reference Manual*, 4.9 ed. CGAL Editorial Board, 2016. [Online]. Available: <http://doc.cgal.org/4.9/Manual/packages.html>



Haichao Miao is a Ph.D. student at the Institute of Computer Graphics and Algorithms, TU Wien and at the Austrian Institute of Technology, Austria. His research interests include visualization of bio nanostructures and medical visualization.



Gabriel Mistelbauer is a postdoctoral researcher at the Otto-von-Guericke University Magdeburg, Germany since 2016. Before, he was a postdoctoral researcher at TU Wien, Austria, where he received a Ph.D. (2013) in computer science and a Master's degree (2010) in visual computing. His research mainly focuses on medical visualization, especially visual analysis of blood vessels and therapy planning.



Alexey Karimov was a project assistant at the TU Wien. He holds a MSc and Ph.D. in computer science. His research interests include topology, probability theory, real-time rendering, scientific and medical visualization.



Amir Alansary is a Ph.D. student in the Department of Computing at Imperial College London, UK. His research interests focus on Machine Learning, Medical Imaging and Computer Vision with specific emphasis on 3D medical image analysis including segmentation, motion compensation and super-resolution.



Alice Davidson is a part time data analyst, involved in many projects at the centre for the developing brain at Kings College London. She focuses her time on fetal volumetric reconstructions and segmentation, as well as conducting eye-tracking assessments for ex fetal participants as part of a follow up study. She is also in a training to be a mental health practitioner at Sussex University.



David F. A. Lloyd is a registrar in Paediatric Cardiology based at the Evelina Childrens Hospital, London. He is currently studying for a Ph.D. in advanced fetal imaging at the Division of Imaging Sciences and Biomedical Engineering, Kings College London.



Mellisa Damodaram is an Obstetrics and Gynaecology trainee with an interest in Maternal and Fetal Medicine. Her Ph.D. was on the use of Fetal MRI in foetuses with growth restriction at Imperial College London.



Lisa Story is currently a King's Open Prize Postdoctoral Research Fellow in the Centre for the Developing Brain St Thomas' Hospital (King's College London) having previously completed her Ph.D. at Imperial College London in the field of Magnetic Resonance Spectroscopy in growth restricted fetuses. She is also a senior obstetric registrar in the London deanery.



Jana Hutter did her undergraduate degrees in applied mathematics and received a Ph.D. from the University of Erlangen-Nuremberg in 2014. Since 2014 she is an associate researcher at the Centre for the Developing Brain at King's College London. Her interests include MRI physics, sequence development, diffusion MRI as well as fetal/placental and neonatal imaging.



Joseph V. Hajnal is Professor of Imaging Sciences at Kings College London. He trained as a physicist at Bristol University, England, UK and obtained a Ph.D. in the physics of electromagnetic waves before working in Australia at Melbourne University and the ANU on interactions between atomic beams and laser light. In 1990 he began research in medical imaging. His research group focuses on in vivo imaging, particularly Magnetic Resonance Imaging (MRI). His current research interests include MR data acquisition and processing, image registration and data fusion as well as novel scanner technology, parallel imaging and motion artefact correction. He has invented and pioneered techniques that are now widely used in the medical imaging industry. With Daniel Rueckert, he co-founded a IXICO and also developed the critical motion-tolerant imaging methods that underpin image acquisition in this programme.



Mary Rutherford is Professor of Perinatal Imaging at Kings College London, and is based within the Centre for the Developing Brain at St Thomas Hospital. She is also an Honorary Professor at Stellenbosch University. She originally trained as a paediatrician, specialised in neonatal neurology and became involved in MRI in the late 1980s. She has since worked as a clinical researcher in the field of MRI. Initially studying acquired injury in the term and preterm neonatal brain and since 2002 assessing congenital and acquired abnormalities using fetal MRI. More recently she has started working on imaging the placenta, being PI on an NIH funded project to study placental structure and function with MRI.



Bernhard Preim heads the Visualization Group at the University of Magdeburg (UoM), Germany since 2003. Before, he was a staff member at MEVIS Bremen, Germany. He received a Ph.D. at the UoM in 1998 and a habilitation (venia legendi) in 2002 from the University of Bremen.



Bernhard Kainz is Lecturer in the Department of Computing at Imperial College London. He is researching translational high-performance medical data analysis, machine learning, and interactive real-time image processing techniques as member of the Biomedical Image Analysis, BioMedIA Group in the section of Visual Information Processing. He is honorary research fellow at Kings College London, Division of Imaging Sciences and Biomedical Engineering, St. Thomas Hospital London and collaborates intensively with the department of Bioengineering at Imperial.



M. Eduard Gröller is professor at TU Wien, Austria, and adjunct professor of computer science at the University of Bergen, Norway. His research interests include computer graphics, visualization and visual computing. Dr. Gröller became a fellow of the Eurographics association in 2009. He is the recipient of the Eurographics 2015 Outstanding Technical Contributions Award.

Multiscale Visualization and Scale-Adaptive Modification of DNA Nanostructures

Haichao Miao, Elisa De Llano, Johannes Sorger, Yasaman Ahmadi, Tadija Kekic, Tobias Isenberg, M. Eduard Gröller, Ivan Barišić, and Ivan Viola

Abstract—We present an approach to represent DNA nanostructures in varying forms of semantic abstraction, describe ways to smoothly transition between them, and thus create a continuous multiscale visualization and interaction space for applications in DNA nanotechnology. This new way of observing, interacting with, and creating DNA nanostructures enables domain experts to approach their work in any of the semantic abstraction levels, supporting both low-level manipulations and high-level visualization and modifications. Our approach allows them to deal with the increasingly complex DNA objects that they are designing, to improve their features, and to add novel functions in a way that no existing single-scale approach offers today. For this purpose we collaborated with DNA nanotechnology experts to design a set of ten semantic scales. These scales take the DNA's chemical and structural behavior into account and depict it from atoms to the targeted architecture with increasing levels of abstraction. To create coherence between the discrete scales, we seamlessly transition between them in a well-defined manner. We use special encodings to allow experts to estimate the nanoscale object's stability. We also add scale-adaptive interactions that facilitate the intuitive modification of complex structures at multiple scales. We demonstrate the applicability of our approach on an experimental use case. Moreover, feedback from our collaborating domain experts confirmed an increased time efficiency and certainty for analysis and modification tasks on complex DNA structures. Our method thus offers exciting new opportunities with promising applications in medicine and biotechnology.

Index Terms—Nano, nanotechnology, assembly, multiscale, abstraction, DNA, origami, scale-adaptive modification

1 INTRODUCTION

All organisms have their construction plan stored in the exceptional macromolecule DNA. While the data storage capability of DNA using the four bases Adenine (A), Cytosine (C), Thymine (T) and Guanine (G) is broadly recognized, other outstanding features are less known. For instance, DNA double-helix molecules are extremely stable: they have a length of several centimeters but a diameter of just two nanometers (nms). Due to its great chemical stability, DNA can even be recovered from fossils several hundreds of thousands of years old [37]. Moreover, DNA can also fold into chemically reactive 3D structures such as the horseradish peroxidase mimicking G-quadruplex, which is an important signal reporter in biotechnological applications [49]. These properties as well as its sequence-dependent self-assembly capability and modern automated DNA production facilities have led to DNA becoming a building material for complex 2D and 3D nanoscale objects.

The resulting potential applications in medicine and biotechnology have caused DNA nanostructure research to progress rapidly. DNA vessels to transport and release drugs specifically at cancer cells, e. g., have been proposed [3, 16]. Andersen et al. [3] presented a DNA box with a controllable lid that can trap a cargo inside. Even significantly more complex structures such as DNA robots have already been developed [43]. With the increasing complexity of DNA nanostructures, however, novel computational DNA design concepts are needed because the current tools have been developed just for simple, static DNA objects. The current *in silico* design of nanoscale objects is a convoluted process. The user has to manually consider many aspects to create a reliable structure that can be self-assembled *in vitro* later on.

We address this problem with a new multiscale visualization that allows experts to deal with DNA nanostructure objects of increasing complexity, to observe, inspect, and manipulate them at the level best suited for a given task, and to connect the different levels with each other using a continuous semantic abstraction that distributes the complexity among the scales. This approach allows domain experts to break down the DNA building blocks of their nanostructures into different distinguishable parts and thus different descriptions at several levels. Beginning at the lowest level of the atoms with their bonds, our semantic abstraction gradually simplifies the depiction and allows the experts to deal with increasingly complex nanoscale objects and their intended behavior. Moreover, we describe interactions with our representation that are designed in such a way that they are interpreted similarly, regardless of the scale at which they are being used. We implement our approach using a custom data hierarchy for DNA nanostructures within SAMSON [35], an established tool for the adaptive modeling and simulation of nanosystems. In summary, we contribute

- a problem characterization of designing DNA nanostructures and analyzing the problem from the perspective of visualization,
- the enhancement of the nanostructure design process by reducing complexity through a continuous semantic abstraction of the DNA representation across many multiple scales,
- the integration of consecutive detailed and abstracted representations and the seamless, cohesive transition between them,
- scale-adaptive interactions that facilitate consistent manipulations of the nanoscale objects on multiple scales, and thus
- a general visual representation of DNA nanostructures that is able to depict the result according to different DNA design concepts.

2 DNA NANOSTRUCTURES

The creation of nanoscale architecture consists of two steps: (1) the object design on the computer and (2) work in the laboratory to experimentally test the design. Because it can take weeks to months to carry out the experiments and finally assemble the objects *in vitro*, the computational design is crucial in assessing whether the structure will self-assemble *in vitro*. Visualization thus plays an essential role in optimizing an efficient *in silico* design of these structures. Before discussing the related work on molecular visualization and the requirements for our new multiscale visualization, we review the domain background and terminology.

- Haichao Miao, Johannes Sorger, M. Eduard Gröller, and Ivan Viola are with TU Wien, Austria. e-mail: {miao | sorger | groeller | viola}@cg.tuwien.ac.at.
- Haichao Miao, Elisa De Llano, Yasaman Ahmadi, Tadija Kekic, and Ivan Barišić are with the Austrian Institute of Technology. e-mail: {elisa.dellano.fl | yasaman.ahmadi.fl | tadija.kekic.fl | ivan.barisic}@ait.ac.at.
- Tobias Isenberg is with Inria and Université Paris-Saclay, France. e-mail: tobias.isenberg@inria.fr.
- M. Eduard Gröller is also with the VRVis Research Center, Austria.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

2.1 DNA

DNA (Deoxyribonucleic acid) is a macromolecule that consists of two single strands, paired and twisted around each other to form the characteristic double helix structure. The single-stranded DNA consists of smaller units called *nucleotides* or *bases* (A, G, C, T). A nucleotide is composed of a *backbone* and a *side-chain*. The backbone comprises a *sugar* with an attached *nitrogenous base* and a *phosphate group*. Nucleotides are connected via the phosphate group (5' end) with the sugar of the next nucleotide (3' end). Single-stranded DNA can bind to another strand via *hydrogen bonds* at their side-chains using base pair complementarity: T bonds with A and G bonds with C. A sequence of these nucleotides defines the specificity under which two single strands can be paired to create the *double-stranded DNA*, or *duplex*. Researchers in the field of DNA nanotechnology use these outstanding properties to manufacture custom shapes from synthetic DNA.

2.2 DNA Assembly Concepts

In 1982, Seeman [46] proposed to use the base pairing rules to bind parts of DNA strands together. Since then, DNA has become a building material for complex 2D and 3D nanoscale objects due to its sequence-dependent self-assembly capability and the establishment of modern automated DNA production facilities. In the last decade, three different design approaches were developed to create DNA nanostructures: DNA origami, wireframe, and DNA tiles. The general principle, however, is the same—they all are based on folding a DNA strand by making the parts of the sequence complementary to different strands. The specific DNA sequences thus have to be determined beforehand to facilitate the spontaneous assembly. The long strand is therefore often referred to as *scaffold*, while short strands are called *helper* or *staple strands*.

The most popular DNA nanostructure designs are based on a technique referred to as **DNA origami** [41]. It uses hundreds of 20–60 bases helper strands to fold a long DNA scaffold (approx. 7300 bases) into a condensed DNA object. High magnesium concentrations are required to facilitate the folding reaction and to stabilize the DNA origami nanostructures. The software caDNAno [14] supports the design of such objects [17]. In caDNAno, the user draws the schematic target structure in two orthogonal 2D interfaces. The rules of DNA origami are then embedded into the schematic representation. This process facilitates fast prototyping and provides the domain experts, in combination with external online resources such as Cando [28], information on the stability of the structure. Furthermore, the schematic approach of caDNAno always aligns the duplex in parallel in one of the 2D views, whereas the target structure is three-dimensional. This setup requires the user to mentally reconstruct the 3D shape, for example to understand design constraints. To address this problem, caDNAno also provides a 3D interface that depicts the duplexes of the target structure as tubes. Detailed information on its composition, however, are only available in the 2D schematic view. The schematic view, therefore, is appreciated by many researchers for its high flexibility. With progressively intricate designs of nanoscale objects, the schematic view becomes overloaded with many connections and is, therefore, increasingly difficult to read. In our work, we ensure that we can load caDNAno designs and then propose our multiscale visualization as a complementary approach to enhance the current workflow.

In contrast to DNA origami, the so-called **wireframe** design approach allows experts to create nanostructures that are stable at physiological ion concentrations, a characteristic that is important for potential *in vivo* applications. The wireframe structures also comprise a long scaffold and short helper strands but are less condensed and, thus, structurally more flexible. vHelix [8] and Daedalus [29] are the most popular tools to create these DNA objects. Daedalus uses an automatic approach [52] that takes a closed polyhedron as an input and produces the scaffold and staple strands in the atom representation as the result. This process, however, leaves the user with little control and only the high-level target geometry can be modified as an input. The other tool, vHelix [7], defines a pipeline for creating the target geometry using several scripts. vHelix allows the user to interact with the object on a nucleotide level such as breaking bonds or placing new nucleotides. In

contrast to Daedalus, however, less symmetrical structures are created and the pipeline works only semi-automatically.

Both Daedalus and vHelix need to route the scaffold over all edges of a target geometry. This approach is equivalent to finding an Eulerian circuit in a graph whose nodes and edges are the vertices and edges of the target geometry. Moreover, this problem is NP-hard [18] and both tools handle this complexity differently. vHelix might introduce duplicated edges to ensure a solution to the problem and runs an efficient systematic search to find an optimal routing. Daedalus, in contrast, doubles all edges and introduces new vertices to produce a cycle graph where finding an Eulerian circuit is a trivial problem. After the routing is finished, both tools apply different stapling algorithms to complete the DNA nanostructure. The disadvantage of this method is that, once the Eulerian circuit has been calculated and we wish to introduce new edges, the whole computation has to be repeated. We see our proposed method as complementary to these design concepts as we facilitate intuitive manipulations on all levels of detail, which is difficult to achieve with the existing methods.

The third design approach is based on **DNA tiles** that form periodic and large DNA nanostructures such as tubes and sheets. A single tile comprises only a few short single strands and is, in general, designed by hand due to its low complexity.

2.3 Overall In-Silico Design

In all of the mentioned three concepts, the experts start by creating an abstract description of the targeted geometry—a high-level plan of the objects they want to create. They describe the overall shape by vertices, edges, and sizes—without much attention to DNA-specific structural behavior. Next, they route the scaffold strand to create the overall shape and then place staples along the scaffold to fixate the shape. This process creates crossovers at appropriate locations to join parts of the scaffold together. The sequence of a strand also determines the stability and structural behavior [36], both of which have to be considered carefully during the process. For example, the melting temperature is assessed by the experts using external tools to roughly understand the stability conditions of the structure [28]. With this assessment in place, the experts further optimize the structure's stability. Finally, the atomistic details are of interest. Low-level modifications can be carried out on the atomic scale where the nucleotides are modified to add additional function, such as creating binding sites for specific molecules. When a design is finished, the experts return to high-level tasks such as connecting different architectures with each other.

All three design principles described above allow domain experts to create DNA nanostructures that can be assembled *in vitro*. However, they are all based on modeling the DNA from scratch, rather than on effectively modifying existing assemblies. As technology moves toward increasingly intricate structures with many levels of detail, better computational tools are needed to assist in the design process. Optimal tools would support an iterative design process in which experts examine the created structures at many levels of abstraction. This multiscale approach would allow them to assess the global shape and function as well as the mentioned smaller-scale and low-level properties.

The ability to inspect and modify the structure at many different levels of detail is thus essential for the whole process. High-level tasks are tedious to be carried out on a low-level representation, and vice versa. A multiscale approach would allow the experts to effectively assemble and test designs *in silico*, before testing them *in vitro*. For this purpose, we extend the assembly concepts described above with a multiscale visualization approach that supports the inspection and modification of DNA nanostructures in a scale-adaptive way.

3 RELATED WORK

Visual models of biomolecular structures used for the depiction of DNA at atomistic detail can be divided into *atomistic* models and into *illustrative* (visually abstracted) models [20]. *Atomistic* models directly depict the chemical *bonds* or the *surface* of a molecule. Such a realistic depiction is necessary if the atomic structure itself is essential for determining molecular properties. The bonds convey the chemical properties of a molecule and are most commonly represented by (ball

and) stick models. We use this representation in our *Scale 2. Surface models*, in contrast, depict the interface between a molecule and its environment. Kozlikova et al. [25, 26] and Alharbi et al. [2] gave a comprehensive overview of such atomistic representations.

In large-scale models like a complex DNA machinery, large numbers of atoms have to be displayed simultaneously. To facilitate the rendering of such complex scenes at interactive rates, many approaches use instancing and proxy geometry/impostor rendering [6, 31, 32]. In addition, *illustrative* and abstract models are used to convey features that are not or hardly visible in realistic atomistic visual representations. Further, they allow for a sparse, occlusion reducing representation of the dense atomistic data. Especially in large molecule complexes like DNA strands, higher-level molecular properties such as the overall shape and the chemical stability are typically of more interest than the atomistic detail. With the introduction of a cartoon representation for molecules [40], other abstract illustrations emerged, such as ribbons to schematize the helix structure, or arrows to indicate the secondary structures of molecules [22, 53]. These visually abstracted representations are, specifically used for the visualization of DNA (and RNA) structures. The ladder-like double helix of the DNA is typically represented by a ribbon or tube for the phosphate-sugar backbone and by sticks or ellipsoids for the nucleotide bases. We use several such abstractions in the higher scales of our representation.

Ellipsoids [1] and beads/spheres are used as generic abstractions for coarse-grained representations of DNA structures. A related representation was chosen by Benson et al. [7], which relied on spheres and arrows to depict the nucleotides and their direction in the design of DNA wireframe structures. The authors colored the spheres to indicate their affiliation to the respective staple or scaffold strand. We extend this idea and separate the information encoded in the spheres into our *Scales 3–4* (see Fig. 2). Coarse-grained abstractions, in general, significantly reduce the geometric and visual complexity while retaining the high-level shape of a structure [13, 48]. Many of the listed examples are supported as standard representations by a wide range of molecular visualization tools (VMD [23], PyMOL [44], Chimera [15], and Maestro [45]) that are also used by our partners. In contrast to the visually more abstract coarse-grained representations, Cipriano and Gleicher [10] simplified the form directly from the molecular surfaces to preserve significant shape features and use surface glyphs to visually encode additional properties. In another approach, Cipriano et al. [11] used multiscale shape descriptors for surface meshes to statistically characterize surface regions of varying size. Later, Cipriano et al. [12] proposed a web-based tool for the analysis of molecular surfaces. The underlying algorithm abstracts the natural bumpiness of molecular surfaces, while preserving large-scale structural features. Ertl et al. [19] combined isosurfaces of the ion density around molecules with streamlines that indicate ion direction and velocity to analyze the motion of ions around the DNA in a nanopore. In our case, the overall shape of the DNA structure and, consequently, of the double helix is of the highest importance. In *Scale 3*, we abstract the atoms that of individual nucleotides into spheres that correspond to a beads representation. Instead of using molecule shapes, we derive the visual properties of the various scales from the atom positions and the chemical properties.

Additional properties of molecular structures such as the base pair type or the chemical stability can be included in the three-dimensional representation of the molecules [6] or can be conveyed by multiple linked views. Bernier et al. [9], for instance, used a combination of 1D plots, 2D sequence diagrams, and 3D visualizations for the visual analysis of RNA structures in ribosomes. We directly encode the information in the color of the abstracted representations.

Approaches that support multiple molecular representations have also been proposed by several authors. Bajaj et al. [6] revealed molecular properties by blending together several biochemically sensitive level of details of molecular representations. Guo et al. [21] transitioned between LOD representations based on view distance. Van der Zwan et al. [33, 51] seamlessly transformed the visualization of a molecule along three independent abstraction axes: the structural abstraction level, the visual stylization, and the support of spatial perception. Krone et al. [27] proposed a molecular surface extraction

approach that can also be used for LoD renderings. The user can interactively adjust on a continuous scale the display of structural detail between atomic resolution and various levels of surface abstractions. We apply the same principle of seamless transitions between different representations of molecules. However, in our case the different visual scales are specifically designed to support certain tasks and we cover a large range of semantic scales. Asbury et al. [4] offered a visualization of the human genome at multiple scales with the GENOME tool. Similarly to our solution, each scale is intended to support a specific task. In GENOME, however, the different scales are displayed in separate windows while we provide a single interactive view.

4 MULTISCALE DNA NANOSTRUCTURE VISUALIZATION

The overall goal with this work is to address the needs of our collaborators for a better visualization and interaction support for designing DNA nanostructures as outlined in Sect. 2. For this purpose, we started to closely cooperate with a team of domain scientists working on DNA-based molecular diagnostics and DNA nanotechnology. Their research interest focuses on the development of highly innovative detection technologies including the creation of a cell-drilling DNA robot [5].

We started with a planning period over several months that Munzner [34] would describe as the characterization of the domain. During this time, we had six meetings involving junior and senior researchers from the visualization and the DNA nanotechnology fields to brainstorm research challenges and to discuss multiple conceptual ideas to address these challenges. Our main problem here was to find the appropriate visual encoding for the DNA nanostructure design-process. The resulting discussions of different forms of abstraction for the DNA led to the insight that no single visual encoding would suffice to depict the intricate structures, in particular if also taking the interactive modeling aspect into account. We thus began to integrate several semantically different representations into a multiscale visualization.

Subsequently, this paper’s first author (FA) joined the DNA nanotechnology group and spent 2.5 days per week for 3 months with them. Most closely involved in the discussions and the development of the concept were two domain experts, a female PhD student in physics and computational biology (C1, 30 years old) and the male Principal Investigator with a background in molecular biology (C2, 35 years old). In addition, a PhD student in chemistry (C3, female, 31 years old) focusing on DNA origami design and a PhD student in bioinformatics (C4, 28 years old, male) working on atomic details provided additional feedback throughout the development. C1 and C3 had a year of experience, and C4 had less than one year of experience working on DNA nanostructure modeling/design, while C2 had 10 years of experience in single-stranded DNA design. All of them are co-authors of this paper. Occasionally, also other group members were consulted.

The core team consisting of C1–4 and the FA further developed the concept and the integration of the visualization into the software used by the domain experts. The FA conducted a number of informal interviews about the domain experts’ workflow and interaction challenges. Based on this, the FA together with C1–4 developed, in multiple iterations, the concept of using several distinct semantic levels of abstraction of the DNA structures, occasionally discussing the concept with the remaining authors of the paper. The FA reported the progress in regular presentations to the DNA nanotechnology group, where he gathered additional feedback to improve the concept. The implementation was carried out mainly by the FA and C1.

4.1 Overall Approach

With our multiscale visualization we do not aim to show different spatial scales (i. e., sizes) of the DNA macromolecules. Instead, we intend to display different semantic levels of abstraction. Each of these semantic levels is based on or inspired by traditional representations and fulfills a well defined purpose by highlighting specific features of the structure. While some of these scales reduce the level of detail of the depicted structures, others adapt the visually encoded information. Moreover, we order the semantic abstractions and allow the experts to seamlessly transition between them (see Fig. 1)—inspired by the continuous transition between primary and secondary structures for




Fig. 1: The DNA’s atomic details are increasingly abstracted and seamlessly transitioned across ten semantic scales. Based on this multiscale visualization, scientists apply operations such as DNA strand breaking, concatenating, and connecting at different scales.

proteins from van der Zwan et al. [51]. In our approach, however, we not only explore a much higher number of abstraction stages and connect the molecular data to abstract, higher-level geometric structures (e. g., wireframe representations) but also designed a number of interaction techniques that can be applied in a scale-adaptive fashion. This means that the interactions can be applied in a similar way at a number of levels (lower part of Fig. 1) and the specific interpretation is then automatically adjusted to the given level. This combined approach is thus novel not only in the field of visualization but also, for the first time, allows the scientists to inspect important aspects of the structures such as chemical stability, overall shape, and potential connection sites to other structures in a way that would be difficult to assess in a single visual scale. In addition, the different scales offer different entry points for editing the structures as we will discuss later in Sect. 4.4.

In the following detailed discussion we demonstrate our visual encoding on a sample dataset of a cube-shaped wireframe structure (Fig. 2). The cube was generated by our partners using the method described by Veneziano et al. [52]. Although it is a simple geometric primitive, a cube can be used as a basis for creating more intricate structures, as demonstrated by Andersen et al. [3]. The edges are built by two parallel DNA duplexes with 31 base pairs (bps) each, the edge length is 10.5 nm, and the entire cube consists of 1608 nucleotides (nts), 32903 atoms, one scaffold strand, and 20 staples.

4.2 Ten Semantic Levels of Abstraction

We begin the description at the lowest levels of abstraction. Scales 0–2 thus depict the DNA structures at the highest level of detail available, i. e., at atomic resolution. The main purpose of these scales is to facilitate manipulations of nucleotides on an atomic level. Such edits of the atomic structure are typically localized to particular sites to change the function of specific nucleotides. It is important that atomic composition and local and global geometry are easily available to scientists in order to complete these tasks. With input from our partners, we thus created three different visual encodings of the atomic structures that highlight various aspects.

In **Scale 0** we display the atoms and bonds in the *licorice representation* (Fig. 2(a)). The atom type is indicated by color. This common visualization is appreciated by our collaborators for its representation of the atomic bonds that hardly occlude the overall structure. Moreover, it allows them to efficiently spot the molecular geometry, e. g., the hexagonal and pentagonal shapes of the bases.

Scale 1 uses the *sticks representation* of the atoms and bonds (Fig. 2(b)). Here the bonds are thicker, i. e., the cylinders have a radius of 0.3\AA . This scale conveys the structure of the entire molecule better than Scale 0, especially if observed from a greater distance.

Scale 2 shows only the atoms as balls with a radius of 0.35\AA , without bonds (Fig. 2(c)), highlighting the different atoms of the DNA.

Next, in Scales 3 and 4 we represent each *nucleotide* as a sphere with a radius of 1.6\AA that we call *nucleosphere*. By reducing the visual clutter of Scales 0–2 we achieve a cleaner depiction of the overall

geometric structure and a more direct representation of the chemical composition of the DNA strands. The color of each nucleosphere indicates one of the four bases (T: orange; A: green; G: yellow; C: pink). The connection of the backbone is depicted by cylinders between the spheres with radii 0.32\AA . This encoding allows the scientists to analyze the sequence of the DNA strands. Scale 3 and 4 differ from each other in the spatial encoding of the nucleotides, each encoding conveying a different aspect of the helix structure.

We newly developed the spatial encoding of nucleotides in **Scale 3**. Here we highlight the *potential* crossovers between helical DNA strands. These crossovers are exploited in DNA nanotechnology to connect a long strand at different locations. Because crossovers can only appear at locations where the strands come close (i. e., where the double helices turn towards each other), we encode their proximity (Fig. 2(d)). Existing crossovers are displayed by long gray connections between double helices. We position the spheres in this scale at the center of mass of the atoms in the sugar and the phosphor group, with the weight of each atom set to 1—this weight standardization is crucial for the depiction of the helix turns. Taking the actual atomic weight into account would result in an irregular pathway of the single strands.

The nucleosphere layout in **Scale 4** more clearly shows the nucleotide pairs (Fig. 2(e)). We achieve this representation by setting the nucleosphere positions to the centers of the side-chains. The small turns of the double helices make it easy to observe the pathway of the single strand as well as of the duplex. In more densely packed scenes, this compressed spatial encoding also better conveys the strands that are part of the same duplex. Potential crossovers, however, are more difficult to spot. We deliberately omit the explicit depiction of connections between paired nucleotides to avoid an otherwise cluttered representation. Moreover, the pairing information can be read from spatial proximity alone, according to our collaborators.

In most DNA designs, a long scaffold strand is used as discussed in Sect. 2. While the location, length, and sequence of the staple strands can be adapted during the modification of DNA structures, the scaffold strand and its sequence remain constant. In **Scale 5** we thus visually separate the scaffold strand from the staple strands (Fig. 2(f)) to allow the scientists to focus on the composition of the staples. We achieve the explicit visual separation by replacing the nucleosphere representation of Scale 4 with what is called the *single-strand* representation. The single strand is uniformly depicted as a gray tube of radius 1.6\AA .

In Scales 6 and 7 we display the scaffold and staple strands in the *single-strand* representation. We abandon the nucleosphere base information to allow the scientists to focus on the strands’ spatial composition. To visually separate the staple strands from the scaffold, we color them using a set of eight colors. We call this representation *plaiting* due to its resemblance to basket plaiting. Analogous to Scales 3 and 4, Scales 6 and 7 differ in their spatial strand layout.

In **Scale 6**, we route the single strands along the nucleobases (analogous to Scale 4). This representation clearly shows how the scaffold is approximating the geometric shape of the intended target structure




Fig. 2: The ten scales of our multiscale visualization, from high detail (lowest abstraction) to low detail (highest abstraction).

(the cube in our example). Fig. 2(g) shows how the three edges come together at each vertex of the cube and demonstrates that the edges, at each high-level geometry vertex, are held together by one staple strand. Along the high-level edge, one staple holds two parts of the scaffold together, thus creating rigidity. The long straight edges of the staples at crossover locations are a result of this spatial encoding and do not contain any nucleotides, which would be revealed in a lower scale.

Scale 7, in contrast, depicts the strands along the backbone of the DNA (analogous to Scale 3) to indicate the locations of potential crossovers. While Scale 3 depicts the nucleotide details of the strands, it can be difficult to determine the specific strand to which a sequence of nucleotides actually belongs, especially in cluttered structures. The more abstract representation in Scale 7 clearly separates the individual staple strands from each other, enabling the user a stronger focus on potential crossovers between individual strands. In Fig. 2(h) the arrow indicates a potential crossover location of the scaffold.

In **Scale 8** we merge the paired strands of a double helix into a single tube representing the duplex (Fig. 2(i)), placing it along the center of the two paired strands. This highly abstract representation of the double helix allows the scientist to focus on higher-level properties of the design that are otherwise difficult to see. In this scale, we highlight in red the non-paired regions at the vertices of the target structure. These




Fig. 3: Two metrics for assessing stability: (a) the melting temperature (20–75°C) and (b) the Gibbs free energy (0 to –10.000 kcal/mol). Both properties are mapped to a color scale (red: low stability; blue: high stability).

regions are of particular interest because, in wireframe designs, staple strands that span over two adjacent edges are used. At these points, the DNA double helices have to bend and it has been shown that unpaired sequences such as those depicted in red can help bending (a sequence of Ts or *polyT*) [47]. The polyT region can also be observed in Scale 5, but Scale 8 more clearly shows that each vertex of the cube is modeled with three polyT regions. In addition, the target structure’s geometric shape as well as the crossovers and junctions between duplexes become clearer. The zoomed-in view in Fig. 2(i), for instance, shows that the cube edge has three crossovers.

Scale 9, finally, depicts the input geometry from the mesh-based DNA structure generation ([7,52]). It allows the scientists to verify how faithful the structure generation algorithm approximated their input design. The two parallel strands of Fig. 2(i) represent a single edge of the cube in Fig. 2(j). Angles between edges can be clearly inspected at this scale, an aspect that is of great importance to our partners. For example, it is suspected that certain angles have an adverse effect on the structure’s overall stability.

Together, the ten scales are the foundation of our visualization approach. Next we describe a number of specific representations for visualizing local properties, before focusing on the seamless interpolation of the scales and the interactions they facilitate.

4.3 Stability Assessment

One challenge in DNA nanotechnology is to create structures that are stable under many conditions. The *melting temperature* is a metric that is widely employed to analyze the stability of two-paired synthetic DNA strands. It is defined as the temperature at which 50% of the sample has denatured to single-stranded DNA and can be computed using a thermodynamical model [42]. Binding regions with a high melting temperature thus are more stable than those with a lower temperature.

The *Gibbs free energy* [42] is another stability metric. It is related to the melting temperature, is relevant in DNA structure design, and is important for our collaborators. For negative energies, the binding reaction is spontaneous and the duplex state is preferred over the single-stranded state. This means binding regions with negative energy are more stable than regions with a positive one.

We thus included both metrics in our visual encoding to allow the scientists to identify potentially unstable parts of the structure and to interactively optimize them with the operations introduced in Sect. 5. First, we define the binding regions for which the metrics are computed—i. e., those parts of the duplex that are not interrupted by an end or a junction. Then, we map both metrics to a color scale between red and blue and display them at each binding region directly on the scaffold strand to indicate the local stability (Fig. 3). Both metrics can be displayed at those scales where the scaffold is displayed as a single strand, i. e., in Scales 5–7. In Scale 5, the scientists can analyze the nucleotide sequence that is responsible for the stability. Complementary to this analysis, Scale 6 and 7 show the influence of the staple length and binding regions on these properties. We also allow the users to toggle the visibility of the staple strands when activating the depiction of either metric (Fig. 3).




Fig. 4: Seamless transitions between two neighboring scales through linear interpolation.

4.4 Seamless Transitions Between Scales

Beyond the individual scale representations discussed so far, for being able to successfully interpret and effectively work with the semantic scales the scientists have to be able to mentally integrate the depicted information. We assist them in this task with a seamless transition between the scales, which supports mental integrations better than multiple linked views. In our case, the latter would require at least ten independent views, reducing the space available for each individual view. We thus co-register all scales with each other on the data level and then linearly interpolate between the corresponding positions, radii, and colors of the displayed visual elements. This interpolation allows the viewer to understand the relationship between the scales and interpret their composition correctly. Fig. 4 demonstrates the transition between the scales, showing the scales and the results of the interpolation half-way in-between each pair of consecutive scales.

The transitions between the atomic scales (Scales 0–2) are straightforward: the radii of the bonds and atoms are linearly interpolated. Next (Scales 2–3), we grow, relocate, and change the colors of the atoms until they merge into nucleotides. At the same time, we let a tube appear to represent the backbone and, later, the connection between the nucleospheres. As we transit from Scale 3 to Scale 4 we relocate the nucleospheres from their backbone positions to the newly developed side-chain arrangement. To move to the plaiting representation from here, we first grow the connecting backbone tube of the scaffold strand until it hides the nucleospheres (Scale 5), before doing the same for the staples (Scale 6). Next, we relocate the scaffold and staples back to their backbone positions (Scale 7). Then we relocate both and merge them into a single double helix representation depicted in blue (Scale 8). The final transition merges parallel double duplexes into single edges of the target geometry (Scale 9), in particular to allow the scientists to understand the scaffold routing along the edges of the target structure. This transition avoids Schlegel diagrams to project the geometry from 3D to 2D for demonstrating the scaffold routing as done in previous work [7, 52]. In practice, scientists can precisely and intuitively transition between the ten scales using a slider or by direct scale entry and thus relate neighboring scales to each other.

5 SCALE-ADAPTIVE DNA MODIFICATION

In addition to having a visualization that allows the scientists to view, explore, and understand the different abstraction levels, an essential goal of our work is to enable them to make modifications to their designs. Such operations are crucial for increasing a structure’s chemical and structural stability as well as for the assembly of individual structures into more complex architectures or even robots. These interactions, however, are not independent from the multiscale visualization, but have to be applied to a specific or a range of semantic abstraction levels, depending on the given task. Therefore, we created *scale-adaptive modifications* of the DNA structure at all scales where individual strands are visible. With *scale adaptive* we refer to the execution of each operation so that it is adapted to the semantic granularity, i. e., the visual level of abstraction of each scale. The same operation can thus be carried out on atoms, nucleotides, or single strands, but with an impact that is consistent across all scales. Visual elements can be selected across all scales for this purpose. In Fig. 5, for instance, the selection is carried out directly in Scale 4 by simply clicking on a nucleosphere—in contrast to individually selecting all scattered atoms/bonds of a nucleotide.

Below we first describe three basic operations that we currently support. Then, we discuss how these operations are applied in three specific tasks that are essential for the modification of DNA nanostructures.

5.1 Basic Scale-Adaptive Operations

All of our modifications are single-strand operations, which are widely employed in DNA manipulation. They can be carried out on the atomic scales (0–2), the nucleotide scales (3–5), and the single-strand scales (5–7). It does not make sense, however, to carry out single-strand operations on Scales 8–9, as they are visual encodings of higher-level abstractions, such as duplexes and edges. Nevertheless, the results of the operations are visible on all scales. To facilitate the manipulations, we allow the scientists to select visualization elements based on the depicted scale. On the atomic and nucleotide scales, individual atoms, nucleotides, or the bonds between them can be selected. On single-strand scales, a specific position on the strand can be selected which, internally, corresponds to a specific nucleotide within the strand.

Strand Breaking: This operation breaks the DNA backbone by removing the bond between the P and O 3’ end atom in the 3’ end direction, splitting one strand into two. Experts use this modification to cut strands, to shorten them, or to prepare them for crossovers. Depending on the scale, users can break a bond directly between specific atoms (Scales 0–2), between specific nucleotides (Scales 3–5), or between specific locations on a DNA strand (Scales 6–7). On the nucleotide and single-strand scales, we thus remove the bond that corresponds to the 3’ end direction of a selected nucleotide. On the atomic scales, scientists specify the specific atom at a bond location between nucleotides or the bond itself that they want to remove. We propagate the effect of this operation to every scale as shown in Fig. 1 and Fig. 5.

Strand Concatenating: This operation is inverse to strand breaking and is also shown in Fig. 5. Two strands can be merged if their 3’ and 5’ nucleotide ends are direct neighbors. Analogous to strand breaking, for Scales 3–7 we add the bond that corresponds to the 3’ end direction of a selected nucleotide. In Scale 0–2, the scientists have again full control and can select the specific atoms to create a new bond. With such strand concatenating operations scientists can create crossovers and merge several short strands into a single longer one.

Strand Connecting: In contrast to strand concatenation, *strand connecting* joins two single strands by adding a user-specified sequence of nucleotides (Fig. 6). Scientists use this operation to make a longer connection between two strands. The number of nucleotides that are necessary to bridge the gap between disconnected strand ends can be manually adjusted and is estimated. We divide the length between the two selected ends by the approximate radius of a nucleotide. Our collaborators estimate this radius to be 0.35 nm in analogy to B-DNA structural properties, even though we are manipulating single-stranded DNA. We then ask the user to specify the desired nucleotide sequence (see dialog in Fig. 9(b)). We place this sequence on a straight line between the previously selected nucleotides. We intentionally do not pair the newly created single strand with a complementary strand. Therefore, it does not have the characteristic helix shape. Moreover, the number of nucleotides in the sequence considerably influences the conformation of atoms. The atom placement can thus only be accurately calculated by a molecular dynamics simulation, which is outside the scope of this work. In contrast to the other elementary operations, the strand connecting operation also does not make sense for Scales 1 and 2 because it would be too tedious.

5.2 Applications in DNA Structure Modeling

The scale-adaptive basic operations are illustrated with three higher-level operations, frequently carried out by our collaborators for DNA structure modeling.

Strand Merging: Short strands have a negative impact on the struc-



Fig. 5: Breaking & concatenating single strands, in multiple scales. After changes are done in one scale, we propagate the effects to other scales. Breaking a single strand does not break the double strand, thus Scales 8 and 9 are not visually affected.



Fig. 6: Strand connection with a user-specified sequence. Here, atomic scales are out of scope because the atom and bond conformation is subject to molecular dynamics simulations. The higher-level representations in Scales 8 and 9 are also out of scope.




Fig. 7: Strand Merging. The green and violet staple strands are to be merged. First the crossover is removed by breaking the strand. The strand ends are then concatenated. Views in other scales are included manually for demonstration purposes.




Fig. 8: Creating crossovers using scale-adaptive operations: (a) at a potential scaffold crossover location (arrow), the scaffold is broken, (b) first the right strand then (c) the left strand; (d) the newly created ends are concatenated.

ture’s chemical stability. The scientist may merge two short staples to form a longer strand as shown in Fig. 7. The scientist starts by selecting the respective position to interrupt the crossover of a staple (green in Fig. 7(a)). As mentioned above, Scale 6 is well suited for this task because it prominently shows the crossovers. This operation results in two new, separate strands (green and brown in Fig. 7(b)). Next, the scientist applies the strand concatenation operation to connect the previously separate staples. First the brown and the purple strands in our example are connected to yield the longer purple strand in Fig. 7(c). Then the purple and the green strands are connected (see Fig. 7(d)). The result of the operations is a longer single staple, without a crossover. The closeups in Fig. 7 demonstrate how the scale-adaptive operations are propagated to the nucleotide and atomic scales.

Crossover Removal & Creation: The strand merging operation just discussed was demonstrated for the important operation of removing a crossover. In addition to this crossover removal, DNA nanostruc-




Fig. 9: Prolonging staples with a user specified input sequence.

ture design often requires the opposite operation: creating a crossover. Scientists thus start by inspecting the structures for potential crossover locations, which our Scales 3 and 7 support best. If the double helices are aligned in parallel, these locations appear every 10.5 base-pairs at one full helix turn. For example, Fig. 8(a) shows a potential crossover in the scaffold strand (gray), indicated by an arrow. Again using strand breaking, scientists interrupt the previously continuous scaffold (Fig. 8(b) and (c)). Then they re-connect the loose ends using strand concatenating operations to establish the crossover (Fig. 8(d)).

Strand Prolongation: Beyond the local operations, it is sometimes also necessary to add entirely new strands. These are either new *pillars* to increase the rigidity within a DNA structure or binding sites for other molecules that can add new functional behavior to an existing structure. For both tasks, scientists can now use the strand connecting operation we provide. The example in Fig. 9 demonstrates this operation by adding a diagonal to the cube. First, the scientists select two strand ends (arrows at green and red staples in Fig. 9(a)) and then enter the new nucleotide sequence from the 5’ to 3’ direction (Fig. 9(b)) based on the estimated length (20.5 nm) and suggested number of 60 nucleotides. Our strand connecting operation then adds the new strand (Fig. 9(c)), which can also be inspected in other scales (Fig. 9(d) shows Scale 4).

6 IMPLEMENTATION AND PERFORMANCE

We implemented our concept within a nanostructure design suite, which our partners are developing, based on the established SAMSON [35] platform for fast prototyping in computational nanoscience. The implementation can thus be used for other projects with the SAMSON

framework. Specifically, we use SAMSON’s capability for illustrative rendering of geometric primitives. SAMSON itself takes care of the specific GPU-based implementations of the primitives, while we compute the transitions on the CPU. For our data model we directly extend SAMSON’s data graph and its logic to realize the multiscale visualizations and scale-adaptive modifications. We also use the *nthal* package from the *Primer3* software [50,54] for calculating both melting temperature and Gibbs free energy. Based on a loaded project, we automatically create all scale representations. The user interface comprises a 3D view for basic interactions. The representation is precisely controlled through a single slider to depict the current (interpolated) scale. Depending on the scale, the user can select all visible elements and apply the scale-adaptive operations through respective buttons. We tested the implementation on a laptop with an Intel Core i7 CPU and an Nvidia GTX 1060 GPU. For the datasets shown in this paper, we achieve approx. 20 fps including the transitions—fast enough for a smooth exploration and fluid interaction.

7 EXPERIMENTAL RESULTS AND VALIDATION

Our approach applies to structures from different concepts and we make, e. g., no difference between DNA origami or wireframe structures. First, we discussed the feedback on two different designs of a cage received from focus group discussions with our collaborators. Next, we present a case study that demonstrates how different components can be connected with each other, to create a higher order of nanoscale devices. We show the realization using our approach, compare it to the existent approach, and report a tremendous increase in time efficiency.

7.1 Focus Group Discussions

The cage shown in Fig. 10(a)–(c) is a DNA origami structure with an estimated height of 16 nm. We loaded the structure into our system and demonstrated it to C3, the creator of this object. The familiarity with the data allowed her to make a direct comparison to the existing tools she was used to. In contrast to tools such as caDNAno, she immediately appreciated the ability to analyze the structure in 3D. In contrast to the limited 3D views, existing in some established tools, where one “cannot see much,” she particularly appreciated our novel way to show all important details depending on the chosen semantic scale. She noted that the multiscale concept allows her an “all-in-one” solution for the design of structures with modified behavior that would require scientists with different specializations. While the design of the static nanostructure could be done by her using Scales 3–7, another scientist (e. g., C4) can work on the same dataset and modify the atomic details available in the Scales 0–2. The multiscale concept thus provides them with a collaboration possibility that did not exist before. Furthermore, Scale 4 was appreciated for its “straightforward” representation of the single strands and their sequences. Before, this information could only be obtained by combining two views. Although C3 appreciated the color coding of the nucleotides, she would have preferred an optional labeling, especially while still unfamiliar with the software. In Scale 5, she could clearly see and understand the sequence of staples along the paired scaffold due to the 3D depiction. For her work, she stated, Scales 6 and 7 are the most useful ones as they allow her to inspect the staple placements. A limitation of our approach for DNA origami structures is, however, that we cannot display all scales: the duplex representation of Scale 8, which is computed as the center of two paired bases, is not available in the data and Scale 9 is missing because DNA origami structures have no underlying geometry definition.

When asked about modifications, she saw no need to change this structure because the design was already mature. She stated, however, that the abstractions and visual encodings are “simple” and easy to understand—a prerequisite for efficient modifications. C3 also asked to see the computational predictions of the solution shape, but such a representation would require an advanced finite-element analysis, which is beyond the scope of this work. Furthermore, C3 suggested to create an untwisted and parallel representation of two paired single strands, a feature which we plan to add in the future. Overall, C3 highly appreciated our visualizations and asked us to integrate them with her current tools to enhance the nanostructure design workflow.




Fig. 10: A DNA origami (top) and a wireframe structure (bottom).

In a second focus group, we discussed the wireframe cage shown in Fig. 10(d)–(f), created using the approach by Veneziano et al. [52]. This object is larger and more complex than the cube we used as an example throughout the paper. Our partners’ current approach for modifying such structures is to interact on the atomic level to connect staple strands or to change the input to the entirely automatic pipeline. Considering the large number of atoms and bonds, this is a very tedious process that, according to our partners, was up to now only possible in this way. Moreover, any change to the input polyhedron has a global effect in the current approach, resulting in the scaffold being rerouted and the staples being replaced. With our approach, we allow the scientists to modify the structure at all scales, always with local effects. C1 stated that she is very satisfied with this possibility—it allows her to inspect the structure in a “high-level” visualization that goes beyond atomistic representations. C1 is working on the development of DNA algorithms for wireframes, and she mentioned that our visualizations assist her a lot in the process. In general, the edges and vertices of the targeted architecture can be much easier recognized in Scales 3–8 than with atomic details. The routing of the scaffold strand along the edges is an important task, as described in Sect. 2, and can be inspected in Scales 5 and 6. The placement of staples completes the wireframe design. According to the C1, our Scales 6–7 were often used to inspect the placed staples. Finally, the transitions from Scale 9 to 0 allows her to comprehend the link between the input data and the resulting DNA wireframe design. No conventional tool provides this interaction, which drastically simplifies the design process.

7.2 Case Study

To better understand the practical application of our tool, we created a case study for connecting two individually designed components. The concept of such connections is illustrated in Fig. 11(a) where a *nanotube* is connected to a longer *rod*. With existing tools this is a difficult task. Based on the structures modeled in caDNAno, scientists have to identify the correct strands that are candidates for prolongation/connection in a simple 3D viewer. Then, both components cannot be spatially aligned in caDNAno for the connection—the scientists have to mentally integrate both parts and their individual schematic representations. Next, they have to consider the lengths and distances, without being able to estimate these values from the dimensionless schematics. All these aspects make it difficult to decide on the strands to connect and on the lengths of the needed extensions. According to C3, such a task takes several weeks with the existing tools. Moreover, the scientists have to rely on their experience to be confident that the created structure will assemble *in vitro*, making the existing approach temporally and financially expensive and inefficient.

In the case study, we asked C3 to carry out exactly the same task using our system. She started by loading the two datasets to create the multiscale visualizations, which allowed her to quickly inspect the structure. She first inspected the sequence and then the overall placement of the staples and scaffold. The atomic details were not




Fig. 11: Connecting two components via strand prolongation.

of interest for her at this point. Because C3 was not familiar with 3D interactions, the FA assisted her by arranging the two parts to be connected (Fig. 11(b)). C3 particularly appreciated the spatial arrangement in the same 3D view, which allowed her to inspect both structures for finding the best possible location for the connection. Fig. 11(c) shows views of this inspection. Next, C3 identified six staple strands of the rod and the nanotube, which she wanted to connect. She applied our strand prolongation method to add the appropriate sequences from the 5' to 3' direction using the same sequences as in her traditional tools. Because C3 stated that the strand ends are not always easy to see due to the density of the overall structure, we deactivated the visibility of the scaffold (Fig. 11(d)). One way to improve the interaction would be to highlight the strand ends. While this contradicts our general design guideline of reducing the complexity as much as possible for each scale, it would be possible to show this information on-demand. After C3 had completed all six prolongations (Fig. 11(d)) she inspected the connection in a close-up view in Scale 4. C3 stated that the ability to visually inspect the result with our straightforward visualizations gave her great confidence of the resulting stability. Finally, we exported the connected structure and compared it to the result of the old method, verifying that it was identical. The transmission electron microscopy image of this structure is shown in Fig. 11(e). The entire work on this dataset took C3 less than 30 minutes, demonstrating the potential for a much more economic design process than is possible with existing tools. As a result, C3 and her entire team are very satisfied with the new possibilities.

8 CONCLUSION AND OUTLOOK

With this paper we have provided the first effective way of representing DNA nanostructures and for allowing scientists to manipulate them. We have achieved this result based on two fundamental contributions: First, we derived a meaningful *multiscale semantic abstraction* sequence for representing the DNA structures and, second, we introduced the concept of *scale-adaptive manipulation* to facilitate manipulations. The series of semantic abstractions starts at established atom/bond representations via secondary and tertiary structures. These are then

connected seamlessly to the different domain of high-level geometric shapes, which are the building blocks in nanotechnology. Moreover, with the scale-adaptive manipulations we are the first to discuss and solve fundamental interaction issues in abstract visualizations in which the user freely controls the abstraction level.

With our approach we have demonstrated an application of illustrative visualization concepts to a practical domain with realistic datasets (e. g., 264,005 non-hydrogen atoms in Fig. 10(a)) and use cases from the DNA nanotechnology field. We grounded our work by first analyzing the domain and its needs (i. e., a UWP approach [24, 30]) and then by working embedded in the team of our collaborating domain scientists. Throughout the work, we asked for direct feedback and, after completing our new system, discussed its applicability to solve their relevant tasks (VDAR evaluation [24, 30]). With this analysis we showed the applicability of the overall visualization and interaction design, and also demonstrated a drastic increase in time efficiency. The implementation is integrated into the software suite used in DNA nanotechnology and will soon be released as open source. At the time of writing, the collaborators are in the first stages of integrating our system into their work practices and have already completed the first experiments based on it, as documented in Fig. 11(e).

Of course, our approach is not without limitations. The multiscale concept only lays the foundation for basic operations to allow domain scientists to solve essential tasks. In this ongoing project we will continue to work on novel scale-adaptive operations for multiple scales to create advanced multi-component complexes. Multiple levels of semantic abstractions, for example, may leave us with an extreme number of primitives to display in high-detail scales. Currently we are limited to datasets that have less than half a million of atoms. Larger assemblies can no longer be rendered at high-enough refresh rates. We are exploring the use of view-dependent abstraction techniques [31, 38, 39] to later also be able to process datasets with millions of atoms, which is already possible in some DNA nanostructures. Applying our approach to other domains is limited by requiring a sequence of abstraction stages that can be meaningfully interpolated.

The field of nanotechnology is starting to use the outcome from decades of research in computational geometry and graph drawing. The adopted algorithms are just the tip of the iceberg of all the opportunities for scientific cross-pollination between nanotechnology and computer science. For example, while our modeling happens directly in 3D, a 2D modeling is certainly needed when focusing on a certain detail. 2D views eliminate many ambiguities and orientational issues associated with 3D. As a next step we will investigate techniques for structural local flattening of nanostructures for interactive computer aided design.

We see that currently we are dealing with modeling a structure, modeling a function will likely involve the generation of hybrid macromolecules, partially constructed from nucleotides and partially from aminoacids. Such models will require more complex abstraction schemes, but they will also require understanding the needs for a CAD design of functional components.

Currently we are building the most elementary structural components. Soon there will be larger collections of nanomodels available, forming the *screws-and-bolts* of the nanoscale. On a long term, we envision to build and support a digital warehouse with nanocomponents. Designing nanostructures will be similar to producing macroscopic man-made artifacts such as cars and furniture: more complex objects are a composition of pre-made simpler objects.

ACKNOWLEDGMENTS

This project has received funding from the European Unions Horizon 2020 research and innovation programme under grant agreement No 686647. This work was also partially funded under the ILLUSTRARE grant by FWF (I 2953-N31), ANR (ANR-16-CE91-0011-01), and the WWTF (VRG11-010). The TEM samples were prepared / data was recorded at the EM Facility of the Vienna Biocenter Core Facilities GmbH (VBCF), Austria.

REFERENCES

- [1] E. Abrahamsson and S. S. Plotkin. BioVEC: A program for biomolecule visualization with ellipsoidal coarse-graining. *Journal of Molecular Graphics and Modelling*, 28(2):140–145, Sept. 2009. doi: 10.1016/j.jmgm.2009.05.001
- [2] N. Alharbi, M. Alharbi, X. Martinez, M. Krone, M. Baaden, R. S. Laramée, and M. Chavent. Molecular visualization of computational biology data: A survey of surveys. In *EuroVis Short Papers*, pp. 133–137. Eurographics Association, Goslar, Germany, 2017. doi: 10.2312/eurovisshort.20171146
- [3] E. S. Andersen, M. Dong, M. M. Nielsen, K. Jahn, R. Subramani, W. Mamdouh, M. M. Golas, B. Sander, H. Stark, C. L. P. Oliveira, J. S. Pedersen, V. Birkedal, F. Besenbacher, K. V. Gothelf, and J. Kjems. Self-assembly of a nanoscale DNA box with a controllable lid. *Nature*, 459(7243):73–76, May 2009. doi: 10.1038/nature07971
- [4] T. M. Asbury, M. Mitman, J. Tang, and W. J. Zheng. Genome3D: A viewer-model framework for integrating and visualizing multi-scale epigenomic information within a three-dimensional genome. *BMC Bioinformatics*, 11(1):444:1–444:7, 2010. doi: 10.1186/1471-2105-11-444
- [5] Austrian Institute of Technology. MARA – Molecular analytical robotics assays. Website: <http://maraproject.eu/>, 2017. Visited in March 2017.
- [6] C. Bajaj, P. Djeu, V. Siddavanahalli, and A. Thane. TexMol: Interactive visual exploration of large flexible multi-component molecular complexes. In *Proc. Visualization*, pp. 243–250. IEEE Computer Society, Los Alamitos, 2004. doi: 10.1109/VISUAL.2004.103
- [7] E. Benson, A. Mohammed, J. Gardell, S. Masich, E. Czeizler, P. Orponen, and B. Högberg. DNA rendering of polyhedral meshes at the nanoscale. *Nature*, 523(7561):441–444, 2015. doi: 10.1038/nature14586
- [8] E. Benson, A. Mohammed, J. Gardell, S. Masich, E. Czeizler, P. Orponen, and B. Högberg. vHelix – Free-form DNA-nanostructure design. Website: <http://www.vhelix.net/>, 2017. Visited in March 2017.
- [9] C. R. Bernier, A. S. Petrov, C. C. Waterbury, J. Jett, F. Li, L. E. Freil, X. Xiong, L. Wang, B. L. R. Migliozi, E. Hershkovits, Y. Xue, C. Hsiao, J. C. Bowman, S. C. Harvey, M. A. Grover, Z. J. Wartelle, and L. D. Williams. Ribovision suite for visualization and analysis of ribosomes. *Faraday Discussions*, 169:195–207, Nov. 2014. doi: 10.1039/C3FD00126A
- [10] G. Cipriano and M. Gleicher. Molecular surface abstraction. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1608–1615, Nov./Dec. 2007. doi: 10.1109/TVCG.2007.70578
- [11] G. Cipriano, G. N. Phillips, Jr., and M. Gleicher. Multi-scale surface descriptors. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1201–1208, 2009. doi: 10.1109/TVCG.2009.168
- [12] G. Cipriano, G. Wesenberg, T. Grim, G. N. Phillips, Jr., and M. Gleicher. GRAPE: GRaphical Abstracted Protein Explorer. *Nucleic Acids Research*, 38(Web Server issue):W595–W601, 2010. doi: 10.1093/nar/gkq398
- [13] C. Clementi. Coarse-grained models of protein folding: Toy models or predictive tools? *Current Opinion in Structural Biology*, 18(1):10–15, Feb. 2008. doi: 10.1016/j.sbi.2007.10.005
- [14] N. Conway and S. Douglas. caDNAno. Website: <http://cadnano.org/>, 2016. Visited in March 2017.
- [15] G. S. Couch, D. K. Hendrix, and T. E. Ferrin. Nucleic acid visualization with UCSF Chimera. *Nucleic Acids Research*, 34(4):e29:1–e29:5, Jan. 2006. doi: 10.1093/nar/gnj031
- [16] S. M. Douglas, I. Bachelet, and G. M. Church. A logic-gated nanorobot for targeted transport of molecular payloads. *Science*, 335(6070):831–834, Feb. 2012. doi: 10.1126/science.1214081
- [17] S. M. Douglas, A. H. Marblestone, S. Teerapittayanon, A. Vazquez, G. M. Church, and W. M. Shih. Rapid prototyping of 3D DNA-origami shapes with caDNAno. *Nucleic Acids Research*, 37(15):5001–5006, Aug. 2009. doi: 10.1093/nar/gkp436
- [18] J. A. Ellis-Monaghan, A. McDowell, I. Moffatt, and G. Pangborn. DNA origami and the complexity of Eulerian circuits with turning costs. *Natural Computing*, 14(3):491–503, Sept. 2015. doi: 10.1007/s11047-014-9457-2
- [19] T. Ertl, M. Krone, S. Kesselheim, K. Scharnowski, G. Reina, and C. Holm. Visual analysis for space-time aggregation of biomolecular simulations. *Faraday Discussions*, 169:167–178, July 2014. doi: 10.1039/C3FD00156C
- [20] D. S. Goodsell. Atomistic vs. continuous representations in molecular biology. In R. Paton and I. Neilson, eds., *Visual Representations and Interpretations*, pp. 146–155. Springer, London, 1999. doi: 10.1007/978-1-4471-0563-3_15
- [21] D. Guo, J. Nie, M. Liang, Y. Wang, Y. Wang, and Z. Hu. View-dependent level-of-detail abstraction for interactive atomistic visualization of biological structures. *Computers & Graphics*, 52:62–71, 2015. doi: 10.1016/j.cag.2015.06.008
- [22] A. Halm, L. Offen, and D. Fellner. Visualization of complex molecular ribbon structures at interactive rates. In *Proc. IV*, pp. 737–744. IEEE Computer Society, Los Alamitos, 2004. doi: 10.1109/IV.2004.1320224
- [23] W. Humphrey, A. Dalke, and K. Schulten. VMD: Visual molecular dynamics. *Journal of molecular graphics*, 14(1):33–38, Feb. 1996. doi: 10.1016/0263-7855(96)00018-5
- [24] T. Isenberg, P. Isenberg, J. Chen, M. Sedlmair, and T. Möller. A systematic review on the practice of evaluating visualization. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2818–2827, Dec. 2013. doi: 10.1109/TVCG.2013.126
- [25] B. Kozlíková, M. Krone, M. Falk, N. Lindow, M. Baaden, D. Baum, I. Viola, J. Parulek, and H.-C. Hege. Visualization of biomolecular structures: State of the art revisited. *Computer Graphics Forum*, 2017. To appear. doi: 10.1111/cgf.13072
- [26] B. Kozlíková, M. Krone, N. Lindow, M. Falk, M. Baaden, D. Baum, I. Viola, J. Parulek, and H.-C. Hege. Visualization of biomolecular structures: State of the art. In R. Borgo, F. Ganovelli, and I. Viola, eds., *EuroVis State-of-the-Art Reports*, pp. 61–81. Eurographics Association, Goslar, Germany, 2015. doi: 10.2312/eurovisstar.20151112
- [27] M. Krone, J. E. Stone, T. Ertl, and K. Schulten. Fast visualization of Gaussian density surfaces for molecular dynamics and particle system trajectories. In *EuroVis Short Papers*, pp. 67–71. Eurographics Association, Goslar, Germany, 2012. doi: 10.2312/PE/EuroVisShort/EuroVisShort2012/067-071
- [28] Laboratory for Computational Biology and Biophysics at MIT. Cando. Website: <https://cando-dna-origami.org/>, 2015. Visited March, 2017.
- [29] Laboratory for Computational Biology and Biophysics at MIT. Daedalus. Website: <http://daedalus-dna-origami.org/>, 2016. Visited in March 2017.
- [30] H. Lam, E. Bertini, P. Isenberg, C. Plaisant, and S. Carpendale. Empirical studies in information visualization: Seven scenarios. *IEEE Transactions on Visualization and Computer Graphics*, 18(9):1520–1536, Sept. 2012. doi: 10.1109/TVCG.2011.279
- [31] M. Le Muzic, L. Autin, J. Parulek, and I. Viola. cellVIEW: A tool for illustrative and multi-scale rendering of large biomolecular datasets. In *Proc. VCBM*, pp. 61–70. Eurographics Association, Goslar, Germany, 2015. doi: 10.2312/vcbm.20151209
- [32] N. Lindow, D. Baum, and H.-C. Hege. Interactive rendering of materials and biological structures on atomic and nanoscopic scale. *Computer Graphics Forum*, 31(3pt4):1325–1334, June 2012. doi: 10.1111/j.1467-8659.2012.03128.x
- [33] W. Lueks, I. Viola, M. van der Zwan, H. Bekker, and T. Isenberg. Spatially continuous change of abstraction in molecular visualization. In *IEEE BioVis Abstracts*, 2011.
- [34] T. Munzner. A nested model for visualization design and validation. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):921–928, Nov. 2009. doi: 10.1109/TVCG.2009.111
- [35] NANO-D, Inria. SAMSON – Software for adaptive modeling and simulation of nanosystems. Website: <https://samson-connect.net>, 2016. Visited March 2017.
- [36] W. K. Olson, A. A. Gorin, X. J. Lu, L. M. Hock, and V. B. Zhurkin. DNA sequence-dependent deformability deduced from protein-DNA crystal complexes. *Proceedings of the National Academy of Sciences of the United States of America*, 95(19):11163–11168, Sept. 1998. doi: 10.1073/pnas.95.19.11163
- [37] L. Orlando, A. Ginolhac, G. Zhang, D. Froese, A. Albrechtsen, M. Stiller, M. Schubert, E. Cappellini, B. Petersen, I. Moltke, P. L. Johnson, M. Fumagalli, J. T. Vilstrup, M. Raghavan, T. Korneliusen, A. S. Malaspina, J. Vogt, D. Szklarczyk, C. D. Kelstrup, J. Vinther, A. Dolocan, J. Stenderup, A. M. Velazquez, J. Cahill, M. Rasmussen, X. Wang, J. Min, G. D. Zazula, A. Seguin-Orlando, C. Mortensen, K. Magnussen, J. F. Thompson, J. Weinstock, K. Gregersen, K. H. Røed, V. Eisenmann, C. J. Rubin, D. C. Miller, D. F. Antczak, M. F. Bertelsen, S. Brunak, K. A. Al-Rasheid, O. Ryder, L. Andersson, J. Mundy, A. Krogh, M. T. Gilbert, K. Kjær, T. Sicheritz-Ponten, L. J. Jensen, J. V. Olsen, M. Hofreiter, R. Nielsen, B. Shapiro, J. Wang, and E. Willerslev. Recalibrating *Equus* evolution using the genome sequence of an early Middle Pleistocene horse. *Nature*, 499(7456):74–78, July 2013. doi: 10.1038/nature12323
- [38] J. Parulek, D. Jönsson, T. Ropinski, S. Bruckner, A. Ynnerman, and I. Viola. Continuous levels-of-detail and visual abstraction for seamless

- molecular visualization. *Computer Graphics Forum*, 33(6):276–287, May 2014. doi: 10.1111/cgf.12349
- [39] J. Parulek, T. Ropinski, and I. Viola. Seamless abstraction of molecular surfaces. In *Proc. SCCG*, pp. 107–114. ACM, New York, 2013. doi: 10.1145/2508244.2508258
- [40] J. S. Richardson. The anatomy and taxonomy of protein structure. *Advances in Protein Chemistry*, 34:167–339, 1981. doi: 10.1016/S0065-3233(08)60520-3
- [41] P. W. Rothemund. Folding DNA to create nanoscale shapes and patterns. *Nature*, 440(7082):297–302, Mar. 2006. doi: 10.1038/nature04586
- [42] J. SantaLucia and D. Hicks. The thermodynamics of DNA structural motifs. *Annual review of biophysics and biomolecular structure*, 33:415–440, 2004. doi: 10.1146/annurev.biophys.32.110601.141800
- [43] Y. Sato, Y. Hiratsuka, I. Kawamata, S. Murata, and S.-i. M. Nomura. Micrometer-sized molecular robot changes its shape in response to signal molecules. *Science Robotics*, 2(4):eaal3735:1–eaal3735:11, Mar. 2017. doi: 10.1126/scirobotics.aal3735
- [44] Schrödinger, LLC. The PyMOL molecular graphics system, version 1.8. November 2015.
- [45] Schrödinger, LLC. Maestro 11: The completely reimagined all-purpose molecular modeling environment. Website: <https://www.schrodinger.com/maestro/>, 2017. Visited in March 2017.
- [46] N. C. Seeman. Nucleic acid junctions and lattices. *Journal of Theoretical Biology*, 99(2):237–247, Nov. 1982. doi: 10.1016/0022-5193(82)90002-9
- [47] B. Suter, G. Schnappauf, and F. Thoma. Poly(dA·dT) sequences exist as rigid DNA structures in nucleosome-free yeast promoters in vivo. *Nucleic Acids Research*, 28(21):4083–4089, Nov. 2000. doi: 10.1093/nar/28.21.4083
- [48] V. Tozzini. Coarse-grained models for proteins. *Current Opinion in Structural Biology*, 15(2):144–150, Apr. 2005. doi: 10.1016/j.sbi.2005.02.005
- [49] P. Travascio, Y. Li, and D. Sen. DNA-enhanced peroxidase activity of a DNA-aptamer-hemin complex. *Cell Chemical Biology*, 5(9):505–517, Sept. 1998. doi: 10.1016/S1074-5521(98)90006-0
- [50] A. Untergasser, I. Cutcutache, T. Koressaar, J. Ye, B. C. Faircloth, M. Remm, and S. G. Rozen. Primer3—New capabilities and interfaces. *Nucleic Acids Research*, 40(15):e115:1–e115:12, 2012. doi: 10.1093/nar/gks596
- [51] M. van der Zwan, W. Lueks, H. Bekker, and T. Isenberg. Illustrative molecular visualization with continuous abstraction. *Computer Graphics Forum*, 30(3):683–690, May 2011. doi: 10.1111/j.1467-8659.2011.01917.x
- [52] R. Veneziano, S. Ratanalert, K. Zhang, F. Zhang, H. Yan, W. Chiu, and M. Bathe. Designer nanoscale DNA assemblies programmed from the top down. *Science*, 352(6293):1534:1–1534:15, June 2016. doi: 10.1126/science.aaf4388
- [53] M. Wahle and S. Birmanns. GPU-accelerated visualization of protein dynamics in ribbon mode. In *Proc. IS&T/SPIE Electronic Imaging*, pp. 786805:1–786805:12. International Society for Optics and Photonics, 2011. doi: 10.1117/12.872458
- [54] Whitehead Institute for Biomedical Research, S. Rozen, A. Untergasser, M. Remm, T. Koressaar, and H. Skaletsky. Primer3. Website: <http://primer3.sourceforge.net/>, 2016. Visited in March 2017.

DimSUM: Dimension and Scale Unifying Map for Visual Abstraction of DNA Origami Structures

H. Miao,^{1,2} E. De Llano,² T. Isenberg,³ M. E. Gröller,^{1,4} I. Barišić,² and I. Viola¹

¹TU Wien, Austria

²Austrian Institute of Technology, Austria

³Inria and Université Paris-Saclay, France

⁴VRVis Research Center, Austria




Figure 1: A nanotube transitions from a 3D shape model to a diagrammatic 2D representation and finally to a clean 1D alignment of single strands. Orthogonal to the dimension axis is the scale axis, which shows different semantic scales. The orange rectangle tracks a single strand through the dimensions and scales.

Abstract

We present a novel visualization concept for DNA origami structures that integrates a multitude of representations into a *Dimension and Scale Unifying Map (DimSUM)*. This novel abstraction map provides means to analyze, smoothly transition between, and interact with many visual representations of the DNA origami structures in an effective way that was not possible before. DNA origami structures are nanoscale objects, which are challenging to model in silico. In our holistic approach we seamlessly combine three-dimensional realistic shape models, two-dimensional diagrammatic representations, and ordered alignments in one-dimensional arrangements, with semantic transitions across many scales. To navigate through this large, two-dimensional abstraction map we highlight locations that users frequently visit for certain tasks and datasets. Particularly interesting viewpoints can be explicitly saved to optimize the workflow. We have developed *DimSUM* together with domain scientists specialized in DNA nanotechnology. In the paper we discuss our design decisions for both the visualization and the interaction techniques. We demonstrate two practical use cases in which our approach increases the specialists' understanding and improves their effectiveness in the analysis. Finally, we discuss the implications of our concept for the use of controlled abstraction in visualization in general.

CCS Concepts

•**Human-centered computing** → **Scientific visualization**; Visualization theory, concepts and paradigms; •**Applied computing** → Computational biology;

1. Introduction

DNA nanotechnology is a young yet rapidly progressing field that aims to design nanoscale devices by employing DNA as the main building block [See82]. The great chemical stability of DNA and the availability of synthetic DNA production facilities make it a preferred material for intricate nanostructures. The resulting, increasingly complex nanoscale shapes show great potential in medicine

and biotechnology. DNA nanotechnology exploits the base pairing capability of DNA and the synthesis of short *staple* strands to fold the long *scaffold* strand (hence the name *origami*). DNA origami is a widely established method for creating DNA-based nanoscale shapes as introduced by Rothemund [Rot06].

For this application domain we designed and developed an interactive data exploration and modeling approach that relies on a

two-dimensional abstraction space. It spans multiple conceptual scales as well as multiple spatial dimensions for the data representation (Fig. 1). We use it as a navigational map to select the best suitable visual representation for a specific task (Fig. 2). With this concept we go well beyond the approach taken in the domain’s popular software tool *caDNAno* [CD16], which implements functionality of the DNA origami method and enables scientists to design nanoscale shapes in silico. However, advanced computational DNA design concepts are needed, as the current tools are not designed for the ever-increasing complexity of DNA nanostructures. Available tools currently only support the design of 3D shapes on 2D DNA diagrams. As a result, the diagrams exhibit a high degree of clutter and overlapping edges. Although domain scientists employ different representations of these complex structures, they are not able to mentally link between them visually. The resulting cognitive load on their workflow increases the overall duration of designing intricate nanostructures and decreases the confidence in their *in silico* design. Moreover, undiscovered errors in the design can lead to failed experiments, which can take from several weeks to months until the self-assembled structures are inspected in a microscope. As a consequence, these problems in the *in silico* design currently slow down the research progress.

With our 2D abstraction map *DimSUM*, we integrate and visually link the representations of the 3D target structure, the 2D diagrams, as well as 1D arrangements to enable designers to effectively apply manipulations in the dimension that is most suited. With *DimSUM* we also integrate seamless transitions that allow researchers to mentally link different representations. With this multi-scale and multi-dimensional approach we thus facilitate researchers to gain more confidence in the design and in the end enable them to create increasingly more intricate structures. They can effectively analyze and understand the structural properties of these complex objects, to gain confidence in the *in silico* design before carrying out the experiment in the laboratory itself. Ultimately this significantly reduces the time and costs needed for typical design operations.

Our interactive visualization design is based on a close collaboration with an interdisciplinary team that aims to create a DNA nanorobot, which consists of several complex DNA origami components. We started by analyzing the requirements and the needed visual representations of this domain. Based on this analysis, we extended a previous multi-scale concept [MDLS*18], which only abstracts the 3D structure on one scale-axis. Indirect 2D and 1D representations are widely used in the domain, due to their occlusion free layout, but they are not suited for estimating the actual 3D shape of the structure. Therefore, we integrate the 2D and 1D representations into a comprehensive and novel abstraction space, spanned by a scale and a dimension axis. Using our resulting abstraction map *DimSUM* as a tool to navigate the different representations, we designed a number of additional interaction techniques that support domain scientists in their tasks. We illustrate these practical tasks needed by our collaborators, using several design studies. Our contributions are as follows:

- An integration of novel and existing DNA nanostructure representations into an interactive abstraction map that allows domain scientists to perceive and understand the conceptual space and geometric relationships between established and new DNA nanos-




Figure 2: The *DimSUM* abstraction map is spanned by a dimension and a scale axis and integrates existing and novel representations.

structure representations. We extend the multi-scale approach of Miao et al. [MDLS*18] with an orthogonal dimension axis to utilize the strength of 2D and 1D layouts to analyze, make sense of, and especially modify the complex structures.

- With this design, we advance the discussion of the concept of abstraction in illustrative visualization. We demonstrate that the spatial layouting of the visual elements along the dimensionality is an important axis of abstraction. The axis goes from a realistic shape depiction to a successfully simpler representation of the elements where alignment details are gradually abstracted.
- We integrate a number of specific interaction techniques that support the *DimSUM* design including abstraction axis snapping, guiding (heat-)maps, and saving viewpoints to support both navigation of the abstraction map as well as the analysis of task-specific interaction behaviors.

2. Background: DNA Origami

DNA origami [Rot06] relies on the self-organizing nature of DNA molecules to construct nanoscale objects. The method uses one long (approx. 8000 bases) *single stranded DNA (ssDNA)* and numerous shorter (20–80 bases) synthetic *ssDNA* strands. These shorter *staple* strands bind to specific regions on the long *scaffold* strand to fold it together into a targeted shape. A strand has a directionality and it goes from the 5’ to the 3’ end of the nucleotide sugar ring, which is relevant for the anti-parallel pairing of *ssDNA*. DNA nanotechnology exploits the base pairing capability of DNA, i. e. two strands only pair and form *double stranded DNA (dsDNA)* if the sequence of nucleotides are complementary to each other. The Watson-Crick complements [WC53] describe that *Adenine (A)* only pairs with *Thymine (T)* and *Cytosine (C)* pairs only with *Guanine (G)*. The helices are then held together by crossovers at designated positions, where one *ssDNA* switches from one helix to an adjacent one. Constraining the adjacent helices to potential crossover locations results in a regular pattern of *dsDNA* alignment, which facilitates the creation of stable structures. Based on this, Douglas et al. [DMT*09] proposed two lattices, a honeycomb and a square lattice, for their *caDNAno* tool.

The DNA used in these experiments is typically described by different structural representations as shown in Fig. 3. Scientists thus have to consider each representation carefully when designing

these nanostructures. The primary structure is determined by the sequence of *ssDNA*. It is usually depicted as a string of bases from the 5' to the 3' direction, which motivates the 1D view we use in our representation later-on. The DNA's self-assembly is determined by the sequence of the scaffold and staple strands. The scientists need abstract representations, which allow them to inspect aspects of the individual strands. For example, long staple strands (> 80 bases) are increasingly costly to synthesize. In addition, the occurrence of **C** and **G** on a strand have to be carefully considered as they tend to be *stickier* than other nucleotides. Finally, a high frequency of crossovers per strand could have an adverse effect on the stability of the structure. All of these analysis steps require a separate view on each staple strand, yet in the state-of-the-art-tools like *caDNA* the strands are interwoven in 2D DNA diagrams.

We employ 2D diagrams to show the secondary structure, which describes the base pairing interaction of the *ssDNA*, that forms the *dsDNA*. The 2D view indicates adjacencies of the double helices by placing the double strands underneath each other. This 2D layout for the design of a 3D structure results in many connections between adjacent double strands in 3D that are not neighbors in the 2D diagram, cluttering it with many overlapping edges as the structure grows in complexity (see the example in Fig. 3). The resulting diagram is difficult to understand—in a situation where a correct design is crucial for the success of experiments.

The three-dimensional shape of the target structure, finally, results from the tertiary structure that encodes the spatial 3D model. This motivates our use of an additional 3D representation. The correct folding of the scaffold strand, for example, can only be achieved if the spatial properties of the strands are carefully considered. The exact folding can only be seen by imaging the results of *in vitro* experiments—the best *in silico* predictions of foldings are computationally very expensive and have limitations. An approximate 3D representation of what the nanostructure would look like if it were to fold properly, however, already allows the domain experts to study the properties on the model. Another advantage is that distances can be estimated in a 3D spatial model—a crucial task for the domain scientists. The combination of this model with already tested design techniques increases the confidence in the *in silico* constructions.

3. Related Work

Our work builds on concepts and methods in DNA nanostructure visualization, visual abstraction, and HCI, which we review below.

DNA nanostructure visualization: For our visualization concept, we integrate existing as well as several novel representations. Among the former, diagrams have been used by domain researchers as simplified DNA representations for a long time. In Rothemund's work on DNA origami [Rot06], for instance, *dsDNA* is depicted in an unrolled way, showing it as colored arrows to indicate the *ssDNA* and their crossovers. Douglas et al. [DMT*09] proposed a tool to rapidly prototype DNA origami structures based on a similar representation. In addition, he depicts the arrangement of the *dsDNA* on a lattice that shows the structure from an orthogonal direction. Based on the regular appearance of potential crossover positions on the *dsDNA*, this representation leads to a hexagonal honeycomb lattice and a square lattice. In addition, a view of the target 3D form

helps the user to understand the final shape of the structure. Using an alternative representation, Benson et al. [BMG*15] described a technique to semi-automatically model polyhedral nanostructures in 3D space, representing nucleotides as spheres. For both types of representations the authors have made well-motivated design decisions that allow the users to build the desired target shape. Various structural details are not described, which are necessary for advanced tasks such as structural modifications of functional nanostructures. The 2D diagrams, for example, are not well suited for tasks in which spatial features have to be considered. Nonetheless, the 2D diagrams by Douglas et al. [DMT*09] are widely known in the field of DNA nanotechnology. We thus integrate them as part of our 2D view, and extend them with seamless transitions to other representations of different dimensionality such as 3D views.

Visual abstraction: The visual abstraction of spatial data is a core concept of illustrative visualization [VI18] that has been applied in numerous domains.[†] For example, the schematization of network data such as streets has been studied in detail, both for traditional geography applications [CdBvK05] and for artistic purposes [Ise13]. A continuous transition of different presentation styles of city models based on task, camera view, and image resolution has been proposed by Semmo et al. [STKD12]. Researchers have also investigated the non-uniformly controlled adjustment of abstraction [AS01] to support, in particular, navigation. Similar to the linear structures in cartography, the DNA strands in our application domain exhibit information that is not always relevant. As an example, for certain analysis tasks the twist of the double helix or the pathway might be irrelevant. In our 1D view we thus straighten the DNA strands and indicate crossover locations—in a way similar to subway lines and maps which share similar characteristics to preserve topological information [Rob12]. For this representation we are also inspired by the straight alignment of single tracks as proposed by Wu et al. [WTLY12] for the purpose of customizing travel paths.

A core advantage of visual abstraction is that it can preserve a visual variable or the (screen) space for the depiction of additional data [VI18]. Cipriano and Gleicher [CG07], for example, abstracted the molecular surface while keeping significant shape features, allowing them to place glyphs on the surface to encode additional properties. We use the same principle in our work to depict data, for instance, about the stability of the DNA origami structures.

A lot more recent work on visual abstraction concerns the dedicated control of the amount of abstraction to support various visualization purposes. One way to control the degree of abstraction is to show different representations of the same object, depending on the distance to the viewer. Parulek et al. [PJR*14] demonstrated this approach also for molecular data, where it serves the visual comprehension and also makes it possible to visualize complex biological environments at interactive rates [LMAPV15]. In addition to such an application, controlled abstraction can also support the exploration of different semantic scales. Miao et al. [MDLS*18], integrated several representations of DNA nanostructures and arranged them

[†] For a more complete survey of related work on visual abstraction we refer to Viola and Isenberg's recent survey/meta paper on the subject [VI18]. Here we only point out aspects that closely relate to or inspire our own work.

on an axis of semantic abstractions that smoothly controls the representation. We use their linear interpolation-based transition and adapt the ten semantic scales to our 2D and 1D layouts. We add the abstraction along different spatial dimensions to facilitate a much more comprehensive exploration of DNA origami structures that also integrates 2D schematics and 1D alignments of DNA strands.

This process leads to the creation of a multi-dimensional abstraction space [VI18] that can be used for the exploration of visual structures by assisting users to mentally integrate different representations. Zwan et al. [vdZLB11] were the first to construct a 3D abstraction space with several components: structural abstraction ('geometric abstraction' in Viola and Isenberg's terminology), abstraction of the illustrativeness, and different degrees of support of spatial perception (the latter two are 'photometric abstractions' in Viola and Isenberg's terms). Recently, Mohammed et al. [MAAB*18] discussed a similar concept for the controlled abstraction of connectomics data. They constructed a 2D abstraction space, the first axis depicting representations of astrocytes and the second axis depicting neurites. In contrast to their work, our approach is more general and not limited to two structures, as our abstraction map is spanned by different axes of abstraction (scale and dimension).

While we employ a similar interaction concept in our work, we augment the novel *DimSUM* abstraction map with specific interaction techniques that specifically support navigation. Furthermore, we enable viewers to control the transition between 1D, 2D, and 3D space without constraints. The dimensional transitions act as dedicated illustration tools to support the understanding of the relationships between the dimensions.

Interaction: The interaction facilitated by the *DimSUM* widget relates to the concept of **multiple coordinated views (MCVs)** [BC87, BMMS91, Rob07, Wil08]: Users interact with the presented data and visually inspect them through different views that are linked to each other. However, our domain users require operations on many scales, dimensions, and their intermediate transition states. Having many views using **MCVs** would not be feasible, as the single views would get too small. The *DimSUM* widget provides a different interaction concept as it allows users to transition between representations shown in another view. Our interaction is thus more akin to approaches like Jianu et al.'s [JDL09] exploration of three-dimensional brain fiber tracts using a previously generated abstract 2D representation. We use animations to transition between the differently abstracted representations. Animation is shown to facilitate comprehension, learning, and memory communication by Tversky et al. [TMB02] and has been previously applied for molecular abstractions by Sorger et al. [SMR*17]. In addition, we also employ brushing and focusing [BMMS91] for selecting elements and observing their transition through dimensions and scales.

4. Methodology and *DimSUM* Concept

Our work has been motivated by collaborations with a team of DNA nanotechnology researchers whose ambitious goal is to create a cell-drilling nanorobot. The first author (C1) of this paper, a visualization researcher, has been working for 11 months as a member of the nanotechnology team and works 50% of his time at their lab. The team is led by the Principal Investigator (C2) who develops innovative

concepts of parts of the nanorobot. The existing tools in the DNA nanotechnology domain to support such complex work are reaching their limits. This is the case in terms of functionality and especially in terms of available visual representations for the features of these structures. From early-on, many challenges emerged regarding the visualization and modeling of the involved complex structures. C1 was thus asked to develop a new visualization concept and, ultimately, a system required for reaching the goal of creating functional DNA nanorobots using stable DNA origami structures. Based on observations, interviews, and focus group discussions, C1 extracted the necessary context and challenges (Sect. 2). Together with the entire team, he derived the tasks and the design considerations for the visualization system. For example, domain scientists analyze different structural properties using different layouts, as described in Sect. 2. Then, they would like to directly examine and manipulate the structure in the appropriate layout. C1 also extensively collaborated with a physicist (C3), and C1 and C3 were responsible for developing a novel software system required for reaching the goals of the ambitious nanorobot project. In addition, a biochemist (C4) focused on the design of DNA origami structures, while another molecular biologist (C5) was specializing in the creation of novel functional structures. C4 and C5 were also responsible for carrying out the experiments and assembling the designed nanostructures in the laboratory, after the *in silico* design was finished.

Depending on the scientists' subgoals and tasks, they use different representations for modeling, editing, analyzing, and creating conceptual designs. For instance, C2 and C4 heavily relied on 2D diagrams where the structure can be quickly modeled. This work requires a deep understanding of the relationship between the 2D diagrams and the associated 3D model—yet they had no direct way of manipulating the structure in 3D space. C5 depended on having a detailed view on the atomic configuration for changing nucleotide conformations at an atomic level with a 3D spatial model. Since he collaborated with C4 for trying to connect the different components together, he needed a way to understand the newly designed structure in 2D as well. This integrated handling of representations across spatial dimensions is not possible with existing tools. The generation of functional structural motifs required C3 to understand the relationship between the conformation of atoms and the twists of the double helices. For scaffold routing, she did not require a realistic spatial model, but needed to see the entire structure in 2D and without visual occlusion. To evaluate the potential crossovers, she relied on analyzing potential crossover locations in a 3D view and also on being able to manipulate the structure in the same view. After finalizing the design, the experts exported the sequence of single strands as a simple sequence of the bases (Fig. 3). While it was common that this sequence of strands would still be modified, there was no way to relate these changes back to the initial 2D or 3D representations. In summary, the current *in silico* design process is inflexible and complex. New visualization technology that integrates the different representations is urgently needed.

4.1. *DimSUM*: Integration of Dimension and Scale

As *caDNA* is the state-of-the-art tool to create DNA origami structures, we focused on its data representations. We demonstrate our design concept on a prototypical nanotube, consisting of a smaller




Figure 3: We provide many ways of interacting with and seamlessly transition between multiple structural representations using the *DimSUM*. Here we show which parts of the abstraction map depict equivalent content as the state-of-the-art tool *caDNA*no.

shaft and a wider body, as shown in Fig. 4. We used *caDNA*no’s automated staple placement to create it. The scientists did not apply any further modifications, so that we can demonstrate the use of our approach to analyze and improve the structure. The nanotube has a length of 260 Å and a diameter of 50 Å at the shaft and 120 Å at the body. The scaffold is 1003 nucleotides long and the structure is folded by 24 staple strands. In Fig. 3 we compare the views that *caDNA*no provides and show how we integrated equivalent ones into our *DimSUM* abstraction space. Each of the four boxes depicts a view provided in *caDNA*no. We isolate the different spatial dimension (d) of the representations and treat them as one axis of our abstraction space. In addition, we also use different semantic scales (s) of DNA nanoscale structures [MDLS*18] as a second abstraction axis. By arranging the existing representations appropriately and adding the missing ones, the two axes create a unifying, continuous space to access and interpolate between representations (Fig. 2). We first describe the three different spatial dimensions. Afterwards we discuss the transitions between them and interaction mechanisms to assist in the navigation.

3D Representations. We based our three-dimensional representations on Miao et al.’s [MDLS*18] continuous sequence of multiple semantic representations of DNA nanostructures. These representations range from concrete (all-atoms) to abstract (double strands). We extended their concept by more accurately representing the atoms as spheres with van der Waals radii, encoding the 5’ to 3’ direction of *ssDNA* in Scale 5 and 6, and representing the strand in the most abstract scale as a tube with a realistic radius of 10 Å, which the *dsDNA* occupies (Fig. 4).

These 3D representations realistically encode the spatial arrangement of the DNA origami model after self-assembly. It is important to assess the final overall design (i. e., size and space that the structure occupies) but also to estimate the distances between nucleotides. The multiple scales allow the scientists to visualize the structure without distracting clutter of the 2D views. Most importantly, they




Figure 4: Cylindrical nanotube: examples of representations in 3D, i. e., $d = 3.0$ and variable scales s .




Figure 5: The 2D diagrammatic representation and several examples of views are shown where $d = 2$.

are not required to mentally link to the 3D spatial model from a 2D view when applying modifications because they are able to edit directly in 3D space.

2D Representations. Abstract DNA diagrams are widely employed in the domain due to their visual occlusion-free layout. The arrangement in 2D facilitates an overview of the entire structure. In contrast to the existing representations, we use the same semantic scale representations as in 3D space (Fig. 2) so that the scientists can choose the scale most appropriate to their tasks (e. g., Fig. 5).

Equivalent to the 3D representation, the grey tube depicts the scaffold strand and the colored tubes represent staple strands. In an alternating manner, the staple strands are placed underneath or above the scaffold strand with a distance of 5 Å. The vertical distance between the double strands is now 20 Å which guarantees that the complementary strands are spatially close to each other.

With this approach we distribute the complexity of traditional diagrams among several semantic scales. Depending on the scale, we allow the scientists to view detailed features such as atoms and bond conformations in 2D or to simplify the *ssDNA* to colored tubes. In addition, the depiction of complementary strands in parallel, without the helicity of the strands, is automatically provided by our semantic abstraction in Scales 3–7. This feature allows the user to get a quick glance of the sequence of a strand and its complementary counterpart. Except for translations in one plane and zooming into details, the user can thus explore the entire structure without having to change the view angle to account for visual occlusions.

1D Representations. After the 2D representations have already removed parts of the spatial context in favor of schematization, 1D




Figure 6: Examples of representations in 1D, shown at different scales and sorting by (a) strand ID, (b) GC content, and (c) length.

representations only provide a list of straightened single strands. The strands are vertically arranged, aligned side-by-side with a distance of 20 Å. This abstract depiction is motivated by the last step of the in-silico design workflow, in which the sequence of the strands are simply depicted as linear strings as shown in Fig. 3, which allows scientists to analyze their individual properties as well as the overall list of strands needed for synthesis and extraction in order to build the structure in vitro.

With the strands aligned straight, it also becomes possible to sort them by certain properties, as shown in Fig. 6. A straightforward sorting is by *length*. The length of strands is of interest to our domain scientists as it allows them to identify strands with unusual lengths, which cause high costs in the synthesis as described in Sect. 2. The advantage of the DimSUM abstraction map is that it allows the viewer to relate the sorted strands back to the 2D and 3D representations. Furthermore, we enable the sorting of strands by their respective *content of C and G* with respect to the length which indicates the *stickiness* of a single strand. Such stickiness can either be a desired or an undesired property, depending on the purpose of the staple strand and its placement in the design. The melting temperature and Gibbs free energy can be computed for the binding regions of the strands as done by Miao et al. [MDLS*18], using the thermodynamical model proposed by SantaLucia et al. [SH04]. These two properties can also be used for sorting, roughly indicating the stability of the strands based on their secondary structure. Sorting the strands by the accumulated stability thus enables the user to quickly spot the strands that could compromise the self-assembly, as they pose the weakest link in the chain.

4.2. Seamless Transitions Across Dimensions and Scales

To effectively work with the representations, the scientists have to mentally link the representations. To externalize this mental linking, we support control the abstraction and representation process through seamless transitions. We co-register the data and linearly interpolate the positions of elements to guarantee a uniform change between two adjacent representations. In addition, we combine the position-based transformation with the interpolation of shape and color across the scales. There is no rotational or scaling component in the interpolated transformation.

Fig. 1 demonstrates the general transition from 3D via 2D to 1D and from Scale 7 to 1 at the same time, including the intermediate interpolated structures. In Fig. 7 we show a specific transition




Figure 7: After two strands are selected in 1D, we transition along the orange arrow and show examples of the changing representation.

along the orange arrow, going from $d = 1, s = 6$ to $d = 3, s = 2.3$, and depict some of the interpolated representations. This example demonstrates how we transition from an abstract to a concrete representation of the same structure, allowing the scientists to visually connect the involved parts.

Transitioning from 2D to 3D thus enables the viewer to understand how the structure folds together from the schematic display to its 3D shape. Transition from 2D to 1D depicts how the structural design decomposes into the single strands. Combining transitioning across dimension with transitioning across scales enables the scientists to seamlessly move between any two points of the abstraction space, and to take specific paths in this space to intuitively explore the data. The interpolation guarantees a uniform positional change of the elements. Obviously, the current interpolation of positions does not realistically depict the movement and folding of strands, which would require large-scale molecular dynamics simulations. The simplified transition, however, illustrates relationships and allows the viewer to mentally integrate between any representations.

4.3. Cross-Dimension and Scale Highlighting

In order to boost the coherent connection between the different dimensions and scales, we enable the user to highlight structures at any location in the abstraction map. The viewer can either manually select the elements or automatically highlight certain features and then observe how they transition into other representations.

Manual Highlighting. We allow the viewer to mark elements of interest (atoms, nucleotides, single strands, double strands) at any position in the abstraction map, which we highlight (e. g., in yellow). We then color context elements with a visually less salient color (e. g., in gray) to increase the focus on the highlighted elements. Fig. 7 shows how two strands are selected at $d = 1, s = 6$ in the abstraction map and how they are depicted in other locations in the map as the user transitions along the marked path. The transitions allow the user to observe where the particular strands end up in the 2D schematics and the 3D structure. The figure also illustrates the transition across scales. In the detailed atomic view, the atoms that are part of the strands are highlighted as well.

Automated Highlighting. Interesting features that appear re-




Figure 8: Crossovers are highlighted in green. Short segments between crossovers could have an adverse effect on the stability, which is best observed in 1D.

peatedly in the entire structure can also be detected automatically. Crossovers, for instance, are fundamental in DNA nanotechnology (Sect. 2). A relatively high number of crossovers has an adverse effect on the stability of the strand binding, as the segments could have low stickiness. These crossovers can easily be observed in 3D but visual occlusion may make it difficult to gain an overview. This overview is instead provided by the 2D schematics which indicates how well the double helices are held together. In 2D, however, it is difficult to track an individual strand and to see how many crossovers are performed. In contrast, at $d = 1$ in our abstraction map, individual crossovers can be easily observed. We compute crossover locations in the structure and then highlight (e.g., in green) the two nucleotides that form the crossover from one double helix to a neighboring one. An example is shown in Fig. 8. This way we do not deemphasize the context—it is important to keep the contextual information encoded here. The number of crossovers can be simply determined and the length of the segments between crossovers can be visually inspected. Usually, long segments tend to bind better, provided that the sequence is complementary. This sequence complementarity is best seen at $d = 2, s = (3|4)$. In Fig. 8 we marked the brown strand that has relatively short segments of only four nucleotides between crossovers. This configuration could potentially compromise the stability of the entire structure. The arrow points to the pink strand, which has no crossovers and does not contribute to binding adjacent double helices together. The pink strand could thus be a good candidate for connection with another strand. These features are best discovered in our abstraction map at $d = 1$.

4.4. Abstraction-Adaptive Modifications

Structural modifications are an important part in the workflow of the domain scientists. *Deleting, breaking, connecting, and concatenating* are basic operations that allow them to carry out advanced modifications of a structure. We provide such structural modifications at any point in the abstraction map with a well-defined behavior and scope of the respective effect. While Miao et al. [MDLS*18] proposed *Scale-Adaptive Modifications*, we extend these modifications across dimensions, enabling users to modify a structure in 2D and 1D layouts. A particular task can be carried out at the scale where it makes most sense and the system propagates automatically the modification not only across scales but also across dimensions (*abstraction-adaptive*). Next, we demonstrate two tasks for which this operation across dimensions is particularly useful.




Figure 9: The addition of a bridging strand inside the nanotube requires work in different dimensions and scales. Top row: crossovers are broken in 2D. Bottom row: after breaking, the broken strands are connected with a specific sequence.





Figure 10: Deleting of undesired short strands. Left: a sort by length in 1D reveals the short strands. Top row: one of them is selected in 1D and deleted. The same strand is shown in 2D and 3D. Bottom row: representations in all dimensions show that the strand is removed.

In Fig. 9 we demonstrate a task where a bridging strand is added to the inner part of the nanotube. First we depict the nanotube in 3D ($d = 3, s = 7$) and highlight the two crossovers that need to be broken. Due to the visual occlusion in 3D, which we have just described, we select the crossovers in 2D ($d = 2, s = 5.5$) and break them. We then create a new connection of the nucleotides at the broken crossover location and we add a sequence of ten nucleotides (TACGTAGTTT) in-between in 5' to 3' direction.

In Fig. 10 we demonstrate the removal of a particularly short strand that does not contribute to the stability of the structure. According to the domain scientists, these short strands can cause problems during the self-assembly as they have a high probability to bind to regions and would be *in the way* of the intended folding. In addition, they generate costs during the synthesis. To identify short strands, the scientists can now sort the strand by length and quickly see that there are three staple strands of only length three. They can select one strand and depict it in the other dimensions for demonstration purposes. After removal, Fig. 10 shows that the scaffold strand is unpaired at the respective location. However, the brown staple strand can not be extended along the helical axis because crossovers



(a) Frequently visited places on the abstraction map. We marked the places that were used to create the figures in this work.

(b) Snapping to the discrete dimension and scale or current location. The thumbnails show the explicit viewpoints defined by the user.

Figure 11: We augment the *DimSUM* with additional information. The heatmap depicts the accumulated time a user spent at a particular point. The time in milliseconds is mapped to a heatmap with colors representing the intervals $[0; 1000)$, $[1000, 2000)$, $[2000, 3000)$, $[3000, 4000)$, $[4000, +\infty)$.

cannot be formed at locations where the helical twists move away from each other. This can be seen in our 3D representations.

5. Interaction and Navigation in the Abstraction Map

As we provide a large set of representations, we developed a two-dimensional widget that allows the user to easily change representation by moving an icon on the abstraction map. Our abstraction map not only serves as a concept to organize the transitions across semantic scales and dimension, but we use it directly as a navigation widget to access the abstraction space—in a similar way as done previously by Mohammed et al. [MAAB*18]. In contrast to their approach for connectomics, our abstraction map is completely continuous, despite the transitions between spatial 3D representations and abstract 2D and 1D depictions. In addition, we augment the *DimSUM* abstraction map with a heatmap to mark interesting, frequently visited locations, which indicate particularly relevant representations for a given dataset or task. In a collaborative setting, such highlighting allows team members to understand each others interest when working on the same dataset. The heatmap also provides individual users with a way to track their own interaction and data exploration patterns. Finally, it allows the scientists to understand which parts of the abstraction space they have not yet investigated in detail, to look for potentially interesting representations.

The heatmap in Fig. 11a, for example, depicts the places that we visited to create the figures in this paper, highlighting the distinct abstractions that we used. It shows that we not only visited the discrete dimensions and scales but also the space in-between. For the teaser in Fig. 1, for instance, the transition is clearly marked, indicating the frequent use of this space.

We also support the navigation in the abstraction map by allowing users to snap the cursor to either a discrete dimension (with the *d-key*) or a discrete scale (*s-key*) if the cursor lies within a radius of 0.25 units. Alternatively, users can snap the cursor to its current non-integer location (*g-key*). The display of dotted vertical and horizontal lines (Fig. 11b) further assists this interaction. Finally, we allow users to explicitly mark locations by placing *view points* in the abstraction map, enabling them to come back to this location at a later point in time. As shown in Fig. 11b, we add thumbnails to visually mark such locations.

6. Realization / Implementation

Our system has to effectively integrate all the data required for the visualization in the various dimensions and scales. We employ a model that hierarchically represents the *strands*, the *base pairs*, the *nucleotides*, and *atomic details*, which we describe next. We initialize this model with the original caDNAno 2D diagrams, from which we derive all the information, such as 2D and 3D positions.


CaDNAno diagrams are based on a lattice, which covers a plane. As design principle it is assumed that the helical axis of all ds-DNA are perpendicular to this plane. On the lattice the positions of the *virtual helices* are defined. Double strands can be formed by conducting the scaffold strand through them. In the caDNAno file format it is specified which virtual helices are in fact used in the design by indicating their rows and columns in the lattice. The positions occupied by scaffold nucleotides and/or staple nucleotides within every virtual helix and the routing of every strand across the virtual helices is also specified. The used lattice and the sequence of the scaffold must be provided by the user as it is not contained in the data. For the datasets in this paper we take the standard m13mp18 bacteriophage DNA as scaffold sequence as this is used by our domain scientists. However, any other sequence in the FASTA-format can be input into our system.

The 3D, 2D, and 1D layouts depend on the positions of nucleotides. Each nucleotide defines a position in each dimension. From caDNAno we derive the 2D conformations. We then determine the corresponding 3D positions using the 2D conformation and the lattice information. We can, finally, compute the corresponding 1D locations from a side-by-side placement of single strands. We use the sorting criteria described in Sect. 4 as the different ssDNA are now completely described according to our model. In all dimensions we make use of the geometric properties of a typical double helix to achieve models as realistic as possible. Furthermore, to generate an all-atom 3D model we use idealized base pairs [LO03]. We generate additional scales by aggregating data and, for the transition along the scale-axis, we linearly interpolate colors, shape parameters, and positions from one discrete scale to the next, both as described by Miao et al. [MDLS*18]. For the transition along the dimension-axis, we interpolate positions. By combining both interpolations we generate a morphing-like transformation between any two points in the abstraction map.

We implemented our system using the *SAMSON framework* [NAN16], a platform for fast prototyping in computational nanoscience. We used its capability of rendering geometric primitives such as spheres and cylinders for depicting the primitives in our visualizations. We implemented the user interface using the *Qt framework* [The17]. For calculating the melting temperature and Gibbs free energy, we used the *ntthal* package from the *Primer3* software [WRU*16, UCK*12]. We tested our system on a laptop with an Intel Core i7 CPU and an Nvidia GTX 1060 GPU. For the datasets shown in this work, we achieve 20 fps or more.

7. Results and Discussion

Because we realized our work in close collaboration with domain experts (Sect. 4), we could directly derive their requirements by observing their workflow. The case studies, which we describe next,



(a) Nanocage. A staple needs to be elongated towards the target in the center.



(b) Two candidates for elongation. On the left, depicted in 2D, which does not provide means to estimate the helical twist and positions of strand turns away. On the right, the twists and positions of strands.

(c) The same strands does not provide means to estimate the helical twist and positions of strands.

(d) The logged abstraction map depicts which locations expert C4 visualized in order to analyze the structure and carry out this task.

Figure 12: Demonstration of the detection of specific surface strand.

are the result of focus group discussions as well as of feedback we received during the regular progress reports to the team. We further developed the case studies in two final sessions in which we provided the domain scientists with our system, observed their interaction, and logged their behavior in the *DimSUM* abstraction map. The results demonstrate the effectiveness and efficiency of our approach. Furthermore, the experts could gain insights about their work that were not possible before, in particular with respect to the missing connection between 3D and 2D representation.

7.1. Surface Strand Analysis

A main challenge in the design of nanostructures is that scientists need to rely on realistic estimations of distances and location of the structural elements. Fig. 12 depicts a cage that we use to demonstrate the case. Here, a common task is to identify the strand endings that are on the inside of the cage. These endings are candidates for elongation. Such identification tasks are challenging when designing and modifying DNA nanostructures. With the existing tools such as *caDNAo* it is difficult to determine the strand endings according to expert C2: Not only their locations but also their directions of the helical turn have to be carefully considered.

Therefore, expert C2 suggested to carry out this task in our system. We asked expert C4, who is very experienced with the traditional tool, to find the surface strands and provide us with feedback. She investigated the DNA origami cage with dimensions of $200 \text{ \AA} \times 200 \text{ \AA}$ a dataset with which she was not familiar before. It contains 56 staple strands and one scaffold strand with 2 197 nucleotides, aligned in a honeycomb lattice. Her goal was to identify those strand endings that could potentially be elongated to create a connection towards a target in the center of the cage, as illustrated in Fig. 12a.

While she was familiar with our project, she did not actively

use our system before. We first gave her an introduction into the user interactions and explained the visualization concepts. Then she used our system to explore the data and become familiar with it. We logged her interactions during task executing on the *DimSUM* abstraction map, and we show the resulting heatmap in Fig. 12d. As it is apparent in the figure, the 3D representations at the higher scales were the most interesting ones for identifying the surface strands. According to the expert, she was able to have a straightforward view of the spatial model in 3D and to relate it back to the familiar 2D view. She stated that one big advantage of our approach is that she could get an overview of all possible strand endings and quickly identify those on the inside of the cage. She quickly found two candidates, which we highlight in Fig. 12a. We give a zoomed-in image in Fig. 12b, which shows that these endings are located on the bottom of the lower double helices. She realized that, although both strands end at the bottom, only the left yellow strand turns towards the target in the center of the cage, whereas the right orange one turns away and is, hence, not an ideal candidate for elongation. According to C4, this precise analysis in the planning stage enables her to conduct a time-efficient and accurate *in silico* design and thus reduces the probability of introducing errors that compromise the successful assembly of the nanostructure in the wet lab later. Finally, she transitioned to the 2D representation as shown in Fig. 12c, in order to compare the 2D depictions, which she is familiar with.

This case demonstrates how our approach can assist the domain scientists in tasks that require a good understanding of both the schematic design and the resulting spatial layout. This makes the workflow in such cases much simpler. We also learned that navigating through the 3D scene can pose a challenge to some users such as C4, as they might need prolonged time to familiarize with 3D navigation. Later we discussed this case with C2 and he suggested to automatically detect the surface endings in the future and then to highlight them in our visualization. According to C2, however, our system allowed the scientists to perform the manual identification very well in comparison with existing tools, where the detection of these surface strands is a tedious task.

7.2. Parametrized Generation of Structural Motifs

One domain scientist (C3) has been using our system intensively throughout the development and has embedded it into her workflow, starting at an early point in time. She specializes on the computational design of DNA nanostructures and her goal is to develop methods for the parametrized generation of structural motifs and the computational design of functional nanostructures. The reliance on visualization and the exploratory nature of her work, however, requires her to analyze her results with different methods. She provided regular input and was therefore actively involved in the conceptual development of this work.

In a final interview we asked her again to explain how the system supports her and logged her interaction with our system on her own machine. She loaded the robot/man dataset, used in the work of Castro et al. [CKK*11], to demonstrate how our new approach helped in her tasks. As shown in Fig. 13a, she first selected the shoulder of the man and then transitioned from 3D to 2D to observe where this part of the structure is located in the 2D diagram. She stated that the 3D spatial model is necessary for her to understand the 2D schematics,




Figure 13: Using transition to mentally link between 3D to 2D.

which is sometimes challenging due to the loss of spatial context. According to her, this problem is solved because she can now relate between these two representations in DimSUM. In the heatmap of the logged session (Fig. 13d) we can observe the corresponding frequent transitions. As she is already experienced with our system, she could quickly navigate to the appropriate representation using the proposed interaction techniques. Her feedback regarding transitions between dimensions focused on the development of new algorithms for DNA origami design. The visualization of these transitions made easier for her to understand the DNA origami method in general. She also noted that it can greatly reduce the amount of time spent on debugging DNA nanostructure designs.

Overall, she described that the ability of visualizing the same structure in many dimensional and semantic representations and having the ability to relate between these representations helped her to visually examine the results to validate the correctness of the computational methods she developed. She stated that, for example, a faulty *in silico* design will produce incorrect foldings in the laboratory. Using our system and being able to visually examine the *in silico* design in the proposed manner, she could find errors in the computational methods. She stated that this saves costs and also time spent for assembling potentially erroneous structures.

8. Conclusion and Future Work

In this paper, we discuss the application of visual abstraction to DNA nanotechnology. DimSUM, our two-axes abstraction space goes beyond the work of Mohammed et al. [MAAB*18], which actually encodes geometry and dimensionality on a single axis. We advance the concept of abstraction by demonstrating that dimensionality is an important axis, where different layouts provide specific advantages in analyzing and modifying complex data.

Because we integrated the visual abstractions into an interactive widget we facilitate seamless transitions. This possibility to mentally integrate all representations in our abstraction space, allows domain scientists to include our visualization concept into their workflow, and to improve it. Another contribution includes the use of the DimSUM abstraction map as interaction widget. Traces of user interaction can be taken to optimize visualization settings for particularly focused tasks that do not use the entire visual abstraction space. We could also apply user traces in a systematic design of custom MCV systems. Only a task-dependent subset of views that

have been visited during the interaction will then form the simplified visual interface. A MCV could be advantageous if it allows the user to immediately understand the implications of a modification in views other than the one he or she is currently operating with.

Our DimSUM approach can potentially be generalized to data based on linear structures. The key idea is to share the complexity of data among several layouts and scales, each representation highlighting a certain aspect of the data. As noted in Sect. 3, abstractions of linear structures are important, e. g., in cartography. Here, line representation of streets are continuously generalized depending on the distance to the viewer. Blood vessels also share traits with DNA w.r.t. their different semantic representations (volume, surface, centerlines, etc.) and are typically abstracted to 2D depictions. Another example are fibrous structures in material sciences, where defective fibers could be singled out easily in a 1D arrangement. For the future, it could be of great interest to explore the possibility of localized glimpses into other dimensions and scales while keeping the context—similar to the magic lens [BSP*93] metaphor.

Furthermore, we showed how our visualization concept of interactive abstraction is helpful to the domain of DNA nanotechnology. We provide domain scientists with a proof-of-concept implementation that is already in the daily workflow of some experts, (e. g., C3). One of the limitations of our approach is that we currently are only implementing model modifications. We do not yet provide from-scratch modeling of new designs. Another drawback is that we are not using the positional data available from simulations. DNA origami structures could exhibit twists and bends in their shape model, for example. We want to incorporate these structural predictions in the future to provide better approximations. A more general limitation is that large datasets are increasingly difficult to inspect: elements can become too small for an effective work if viewed from far away. This scalability issue is not necessarily a drawback specific to our approach, but of DNA diagrams in general. Furthermore, users have remarked that the transitions from 2D to 1D are not easy to follow when working with larger datasets. Nonetheless, this concern could be addressed in future work that looks at the design approaches in DNA origami in general to provide a scalable and efficient workflow. Our next research opportunity is to extend the presented work to general DNA nanostructures, including, but not limited to, wireframe structures [VRZ*16, BMG*15], which exhibit a great potential for applications under physiological conditions. In the near future, the implementation of our approach will be released as part of a open-source software toolkit.

Acknowledgments

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 686647. This work was also partially funded under the ILLUSTRARE grant by FWF (I 2953-N31), ANR (ANR-16-CE91-0011-01), and the WWTF (VRG11-010). This paper was partly written in collaboration with the VRVis Competence Center. VRVis is funded by BMVIT, BMVFW, Styria, SFG and Vienna Business Agency in the scope of COMET - Competence Centers for Excellent Technologies (854174) which is managed by FFG. We thank Yasaman Ahmadi and Tadija Kekic for providing valuable input to this work and also David Kouril who created the submission video.

References

- [AS01] AGRAWALA M., STOLTE C.: Rendering effective route maps: Improving usability through generalization. In *Proc. SIGGRAPH* (2001), ACM, New York, pp. 241–249. doi: 10.1145/383259.383286 3
- [BC87] BECKER R. A., CLEVELAND W. S.: Brushing scatterplots. *Technometrics* 29, 2 (May 1987), 127–142. doi: 10.1080/00401706.1987.10488204 4
- [BMG*15] BENSON E., MOHAMMED A., GARDELL J., MASICH S., CZEIZLER E., ORPONEN P., HÖGBERG B.: DNA rendering of polyhedral meshes at the nanoscale. *Nature* 523, 7561 (2015), 441–444. doi: 10.1038/nature14586 3, 10
- [BMS91] BUJA A., McDONALD J. A., MICHALAK J., STUETZLE W.: Interactive data visualization using focusing and linking. In *Proc. Visualization* (1991), IEEE Computer Society, Los Alamitos, pp. 156–163. doi: 10.1109/MSUAL.1991.175794 4
- [BSP*93] BIER E. A., STONE M. C., PIER K., BUXTON W., DE ROSE T. D.: Toolglass and magic lenses: The see-through interface. In *Proc. SIGGRAPH* (1993), ACM, New York, pp. 73–80. doi: 10.1145/166117.166126 10
- [CD16] CONWAY N., DOUGLAS S.: caDNano. Website: <http://cadnano.org/>, 2016. Visited in March 2017. 2
- [CdBvK05] CABELLO S., DE BERG M., VAN KREVELD M.: Schematization of networks. *Computational Geometry* 30, 3 (Mar. 2005), 223–238. doi: 10.1016/j.comgeo.2004.11.002 3
- [CG07] CIPRIANO G., GLEICHER M.: Molecular surface abstraction. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (Nov./Dec. 2007), 1608–1615. doi: 10.1109/TVCG.2007.70578 3
- [CKK*11] CASTRO C. E., KILCHHERR F., KIM D. N., SHIAO E. L., WAUER T., WORTMANN P., BATHE M., DIETZ H.: A primer to scaffolded DNA origami. *Nature Methods* 8, 3 (Feb. 2011), 221–229. doi: 10.1038/nmeth.1570 9
- [DMT*09] DOUGLAS S. M., MARBLESTONE A. H., TEERAPITAYANON S., VAZQUEZ A., CHURCH G. M., SHIH W. M.: Rapid prototyping of 3D DNA-origami shapes with caDNano. *Nucleic Acids Research* 37, 15 (Aug. 2009), 5001–5006. doi: 10.1093/nar/gkp436 2, 3
- [Ise13] ISENBERG T.: Visual abstraction and stylisation of maps. *The Cartographic Journal* 50, 1 (Feb. 2013), 8–18. doi: 10.1179/1743277412Y.000000007 3
- [JDL09] JIANU R., DEMIRALP C., LAIDLAW D.: Exploring 3D DTI fiber tracts with linked 2D representations. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (Nov. 2009), 1449–1456. doi: 10.1109/TVCG.2009.141 4
- [LMPV15] LE MUZIC M., AUTIN L., PARULEK J., VIOLA I.: celVIEW: A tool for illustrative and multi-scale rendering of large biomolecular datasets. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (2015), Eurographics Association, Goslar, Germany, pp. 61–70. doi: 10.2312/vcbm.20151209 3
- [LO03] LU X., OLSON W. K.: 3DNA: A software package for the analysis, rebuilding and visualization of three-dimensional nucleic acid structures. *Nucleic Acids Research* 31, 17 (Sept. 2003), 5108–5121. doi: 10.1093/nar/gkg680 8
- [MAAB*18] MOHAMMED H., AL-AWAMI A. K., BEYER J., CALI C., MAGISTRETTI P., PFISTER H., HADWIGER M.: Abstractocyte: A visual tool for exploring nanoscale astroglial cells. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 853–861. doi: 10.1109/TVCG.2017.2744278 4, 8, 10
- [MDLS*18] MIAO H., DE LLANO E., SORGER J., AHMADI Y., KEKIC T., ISENBERG T., GRÖLLER M. E., BARIŠIĆ I., VIOLA I.: Multiscale visualization and scale-adaptive modification of DNA nanostructures. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (Jan. 2018), 1014–1024. doi: 10.1109/TVCG.2017.2743981 2, 3, 5, 6, 7, 8
- [NANI6] NANO-D, INRIA: SAMSON – Software for adaptive modeling and simulation of nanosystems. Website: <https://samson-connect.net>, 2016. Visited December 2017. 8
- [PJR*14] PARULEK J., JÖNSSON D., ROPINSKI T., BRUCKNER S., YNERMAN A., VIOLA I.: Continuous levels-of-detail and visual abstraction for seamless molecular visualization. *Computer Graphics Forum* 33, 6 (Sept. 2014), 276–287. doi: 10.1111/cgf.12349 3
- [Rob07] ROBERTS J. C.: State of the art: Coordinated multiple views in exploratory visualization. In *Proc. CMV* (2007), IEEE Computer Society, Los Alamitos, pp. 61–71. doi: 10.1109/CMV.2007.20 4
- [Rob12] ROBERTS M. J.: *Underground Maps Unravalled: Explorations in Information Design*. Maxwell J. Roberts, 2012. 3
- [Rot06] ROTHMUND P. W.: Folding DNA to create nanoscale shapes and patterns. *Nature* 440, 7082 (Mar. 2006), 297–302. doi: 10.1038/nature04586 1, 2, 3
- [See82] SEEMAN N. C.: Nucleic acid junctions and lattices. *Journal of Theoretical Biology* 99, 2 (Nov. 1982), 237–247. doi: 10.1016/0022-5193(82)90002-9 1
- [SH04] SANTALUCIA J., HICKS D.: The thermodynamics of DNA structural motifs. *Annual Review of Biophysics and Biomolecular Structure* 33, 1 (2004), 415–440. doi: 10.1146/annurev.biophys.32.110601.141800 6
- [SMR*17] SORGER J., MINDEK P., RAUTEK P., GRÖLLER M. E., JOHNSON G., VIOLA I.: Metamorphers: Storytelling templates for illustrative animated transitions in molecular visualization. In *Proc. SCCG* (2017), Brno University of Technology, Czech Republic, pp. 27–36. 4
- [STKD12] SEMMO A., TRAPP M., KYPRIANIDIS J. E., DÖLLNER J.: Interactive visualization of generalized virtual 3D city models using level-of-abstraction transitions. *Computer Graphics Forum* 31, 3pt1 (June 2012), 885–894. doi: 10.1111/j.1467-8659.2012.03081.x 3
- [The17] THE QT COMPANY: Qt. Website: <https://www.qt.io/>, 2017. Visited December 2017. 8
- [TMB02] TVERSKY B., MORRISON J. B., BETRANCOURT M.: Animation: Can it facilitate? *International Journal of Human-Computer Studies* 57, 4 (2002), 247–262. doi: 10.1006/ijhc.2002.1017 4
- [UCK*12] UNTERGASSER A., CUTCUTACHE I., KORESSAAR T., YE J., FAIRCLOTH B. C., REMM M., ROZEN S. G.: Primer3—New capabilities and interfaces. *Nucleic Acids Research* 40, 15 (2012), e115:1–e115:12. doi: 10.1093/nar/gks596 8
- [vdZLBI11] VAN DER ZWAN M., LUEKS W., BEKKER H., ISENBERG T.: Illustrative molecular visualization with continuous abstraction. *Computer Graphics Forum* 30, 3 (May 2011), 683–690. doi: 10.1111/j.1467-8659.2011.01917.x 4
- [VI18] VIOLA I., ISENBERG T.: Pondering the concept of abstraction in (illustrative) visualization. *IEEE Transactions on Visualization and Computer Graphics* 24 (2018). To appear. doi: 10.1109/TVCG.2017.2747545 3, 4
- [VRZ*16] VENEZIANO R., RATANALERT S., ZHANG K., ZHANG F., YAN H., CHIU W., BATHE M.: Designer nanoscale DNA assemblies programmed from the top down. *Science* 352, 6293 (June 2016), 1534:1–1534:15. doi: 10.1126/science.aaf4388 10
- [WC53] WATSON J., CRICK F. H. C.: Molecular structure of nucleic acids; a structure for deoxyribose nucleic acid. *Nature* 171 (Apr. 1953), 737–738. doi: 10.1038/171737a0 2
- [Wil08] WILLS G.: Linked data views. In *Handbook of Data Visualization*, Chen C.-h., Härdle W., Unwin A., (Eds.). Springer, Berlin/Heidelberg, 2008, ch. II.9, pp. 217–241. doi: 10.1007/978-3-540-33037-0_10 4
- [WRU*16] WHITEHEAD INSTITUTE FOR BIOMEDICAL RESEARCH, ROZEN S., UNTERGASSER A., REMM M., KORESSAAR T., SKALET-SKY H.: Primer3. Website: <http://primer3.sourceforge.net/>, 2016. Visited in March 2017. 8
- [WTLY12] WU H.-Y., TAKAHASHI S., LIN C.-C., YEN H.-C.: Travel-route-centered metro map layout and annotation. *Computer Graphics Forum* 31, 3pt1 (2012), 925–934. doi: 10.1111/j.1467-8659.2012.03085.x 3



Huttengasse 18/3/13
1160 Wien, Austria
+43 676 495 1094
<https://www.cg.tuwien.ac.at/staff/HaichaoMiao.html>
miao@cg.tuwien.ac.at

HAICHAO MIAO

I am a researcher in the field of visual computing with a focus on molecular visualization and modeling. My passion is to collaborate with scientists from exciting fields and to develop novel visual tools.

PERSONAL INFORMATION

Nationality: Austria
Birthdate: 06/04/1987 (in Zhejiang/China)
Languages: German, English, Chinese

EDUCATION

October 2015 – Today
Doctoral program in Engineering Sciences Computer Sciences at the TU Wien, Austria
Advisor: Assoc Prof. Ivan Viola

October 2012 – October 2015
Dipl.-Ing. / MSc. in Medical Informatics at the TU Wien, Austria (graduated with distinction)

October 2007 – June 2012
BSc. in Medical Informatics at the TU Wien, Austria

March 2008 – June 2011
Study of Sinology at the University of Vienna, Austria (not graduated)

RESEARCH EXPERIENCE

Research Associate, October 2015 – now, Institute of Visual Computing & Human-Centered Technology, Vienna University of Technology

Ph.D. Candidate, January 2017 – now, Center for Health & Bioresources, Austrian Institute of Technology

Visiting Research Fellow, July 2018 – August 2018, Visual Computing Group, Harvard University

AWARDS

Best Paper Honorable Mention at the *IEEE Pacific Vis Conference 2017*

Best Paper Honorable Mention at the *Eurographics Workshop on Visual Computing for Biology and Medicine 2015 (EG VCBM)*

Nominated for Distinguished Young Alumnus Award Winter 2015 at the *TU Wien*

Visual Data Storytelling Contest Winner at the *IEEE PacificVis 2019* in Bangkok, Thailand.

PUBLICATIONS:
JOURNAL AND PEER-
REVIEWED
CONFERENCE
PUBLICATIONS

Multiscale Molecular Visualization. [Haichao Miao](#), Tobias Klein, David Kouril, Peter Mindek, Karsten Schatz, Meister Eduard Gröller, Barbora Kozlikova, Tobias Isenberg, Ivan Viola. *Journal of Molecular Biology*, ISSN 0022-2836, 2018.

DimSUM: Dimension and Scale Unifying Maps for Visual Abstraction of DNA Origami Structures. [Haichao Miao](#), Elisa De Llano, Tobias Isenberg, Meister Eduard Gröller, Ivan Barisic, Ivan Viola. *Computer Graphics Forum*, 37(3), 2018.

Multiscale Visualization and Scale-adaptive Modification of DNA Nanostructures. [Haichao Miao](#), Elisa De Llano, Johannes Sorger, Yasaman Ahmadi, Tadija Kekic, Tobias Isenberg, Meister Eduard Gröller, Ivan Barisic, Ivan Viola. *IEEE Transactions on Visualization and Computer Graphics*, 24(1), 2018.

Placenta Maps: In Utero Placental Health Assessment of the Human Fetus. [Haichao Miao](#), Gabriel Mistelbauer, Alexey Karimov, Amir Alansary, Alice Davidson, David F.A. Lloyd, Mellisa Damodaram, Lisa Story, Jana Hutter, Joseph V. Hajnal, Mary Rutherford, Bernhard Preim, Bernhard Kainz, Meister Eduard Gröller. *IEEE Transactions on Visualization and Computer Graphics*, 23(6), 1612-1623, 2017.

Visual Quantification of the Circle of Willis: An Automated Identification and Standardized Representation. [Haichao Miao](#), Gabriel Mistelbauer, Christian Nasel, Meister Eduard Gröller. *Computer Graphics Forum*, 36(6), 393-404, 2017.

CoWRadar: Visual Quantification of the Circle of Willis in Stroke Patients. [Haichao Miao](#), Gabriel Mistelbauer, Christian Nasel, Meister Eduard Gröller. In *Eurographics Workshop on Visual Computing for Biology and Medicine*, 1-10, 2015.

My Erdős number is 3.

PUBLICATIONS:
POSTERS

A Preview to Adenita: Visualization and Modeling of DNA Nanostructures. Elisa De Llano, [Haichao Miao](#), Tobias Isenberg, Meister Eduard Gröller, Ivan Viola, Ivan Barisic. Poster shown at the *3rd Functional DNA Nanotechnology Workshop*. June 6 – 8, 2018.

Interactive Visual Analysis for the Design of DNA Nanostructures. [Haichao Miao](#), Elisa De Llano, Ivan Viola, and Ivan Barisic. Poster shown at the *Nantech Workshop: Nucleic Acid Nanotechnology: From Algorithmic Design to Biochemical Applications*. Espoo, Finland. May 2019

INVITED TALKS

Visualization and Modeling of DNA Nanostructures, September 18th 2018, 1st DNA Engineering Workshop 2018, Plitvice Lakes, Croatia.

Visual Abstraction and Modeling in DNA Nanotechnology, December 14th 2017, King Abdullah University of Science and Technology, Saudi Arabia.

Visualization and Modeling in DNA Nanotechnology, April 2017, Inria, Grenoble, France.

Visualization of Brain Arteries, the Placenta and DNA Nanostructures in the Context of Abstraction, April 2017, State Key Lab - Zhejiang University, China.