



# Enge Integration von Visueller Analyse und 3D Echtzeitgraphik

DISSERTATION

zur Erlangung des akademischen Grades

**Doktor der Technischen Wissenschaften**

eingereicht von

**Thomas Ortner, MMSc.**

Matrikelnummer 00928153

an der Fakultät für Informatik  
der Technischen Universität Wien

Betreuung: Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Eduard Gröller

Diese Dissertation haben begutachtet:

---

Michael Sedlmair

---

Barbora Kozlikova

Wien, 29. Dezember 2020

---

Thomas Ortner



# Tight Integration of Visual Analysis and 3D Real-Time Rendering

DISSERTATION

submitted in partial fulfillment of the requirements for the degree of

**Doktor der Technischen Wissenschaften**

by

**Thomas Ortner, MMSc.**

Registration Number 00928153

to the Faculty of Informatics

at the TU Wien

Advisor: Ao.Univ.Prof. Dipl.-Ing. Dr.techn. Eduard Gröller

The dissertation has been reviewed by:

---

Michael Sedlmair

---

Barbora Kozlikova

Vienna, 29<sup>th</sup> December, 2020

---

Thomas Ortner



# Erklärung zur Verfassung der Arbeit

Thomas Ortner, MMSc.

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, 29. Dezember 2020

---

Thomas Ortner



# Acknowledgements

I want to express my deepest gratitude to my supervisor Eduard Gröller, for guiding me through this journey, and for significantly influencing not only this thesis and the involved research articles, but my research career in general. I would also like to thank Gerd Hesina and Harald Piringer for the original proposal inspiring this work and the opportunity to conduct fundamental research in the scope of the *VISAR* project. And I am honored to have Barbora Kozlikova and Michael Sedlmair as my reviewers.

Secondly, I want to thank my current and former colleagues and friends at the VRVis Research Center for creating the most pleasant work atmosphere, and for their scientific and emotional support. I also want to thank the members of the VisGroup at TU Wien in various configurations for accepting me as one of their own and for integrating me into their scientific- as well as their social gatherings. In particular, I want to thank Johannes Sorger, as my brother in arms in this academic venture, Harald Steinlechner for his ingenious implementations and reminding me when it's time to get a beer, and Michael Schwärzler for his continuous enthusiasm and reminding me when it's time to work on this thesis. I also thank Johanna Schmidt and Thomas Höllt for taking the time to review this thesis and giving valuable feedback. And I am thankful for my long-time collaborators Gerhard Paar from Joanneum Research in Graz, Robert Barnes from the Imperial College of London, and the people from Dibit Messtechnik GmbH in Innsbruck.

Last but not least, I want to thank Madelaine for her continuous support, her fabulous cooking, and her tolerance towards my growing absent-mindedness during the last publication and the final stages of this thesis.

*Large parts of this work have been supported in the scope of the FWF-funded project P24597-N23 (VISAR). All of the presented work was conducted at the VRVis, a COMET K1 competence center, which is funded by BMK, BMDW, Styria, SFG, Tyrol and Vienna Business Agency in the scope of COMET - Competence Centers for Excellent Technologies (854174, 879730) which is managed by FFG. This work was also partly funded by ESA-PRODEX funding, supported by the Austrian Research Promotion Agency under ESA PEA Grants 4000105568 & 4000117520.*





# Kurzfassung

In vielen Fachbereichen, wie zum Beispiel Raumplanung, Bauingenieurwesen, oder Katastrophenmanagement, müssen Expertinnen mit komplexen geometrischen Daten umgehen die mit multivariaten Attributen angereichert sind. Zusätzlich zur visuellen Analyse dieser Attribute, erfordern typische Aufgaben die Lokalisierung und das Verständnis von geometrischen Formen, und das Einschätzen von räumlichen Verhältnissen zwischen geometrischen Objekten und der umgebenden Geometrie, wie zum Beispiel digitale Geländemodelle.

Ein Weg dieses Problemfeld mit Visualisierungs-Design zu adressieren ist der Einsatz von mehreren koordinierten Ansichten, indem man eine geometrische 3D-Ansicht und Attributansichten mittels ‘Brushing & Linking’ kombiniert. Doch wenn man diese Ansichten auf naive Art und Weise koordiniert treten Probleme zu Tage die in der Natur von 3D-Ansichten liegen. Objekte die in Attribut-Ansichten markiert werden, können in der 3D-Ansicht verdeckt sein oder außerhalb des Kamera-Sichtfeldes liegen. Das kann leicht zu Orientierungslosigkeit führen. Aufgaben die auf Lokalisierung, Formverständnis, und dem Einschätzen räumlicher Verhältnisse aufbauen werden erschwert. Letztendlich wird der kontinuierliche Analysezyklus, den koordinierte Ansichten eigentlich ermöglichen sollten, unterbrochen.

In Bezug auf drei Anwendungsgebiete untersuchen wir in dieser Dissertation verschiedene Herangehensweisen um 3D Geometrie- und Attribut-Ansichten visuell zu integrieren. Im ersten Kapitel widmen wir uns dem Bereich der Tunnelinspektion und -dokumentation. Tunnelinspektorinnen sind auf der Suche nach bestimmten Mustern in Tunnelrissdaten um Rückschlüsse auf den Zustand des Tunnels zu ziehen. Wir integrieren eine geometrische 3D-Ansicht und mehrere Attribut-Ansichten zu einer Visualisierungslösung und präsentieren mehrere domänenspezifische Visualisierungs- und Interaktionsstrategien um die zuvor erwähnten Probleme zu kompensieren. Wir schließen das Kapitel mit einem methodologischen Framework das Richtlinien für Visualisierungsdesignerinnen bietet zur Integration von Geometrie- und Attribut-Ansichten. Es besteht aus den Komponenten ‘Geführte Navigation’, ‘Verbesserte Geometrische Darstellung’ und ‘Ähnlichkeitsbasierte Analyse’.

Im zweiten Kapitel erkunden wir die potenzielle visuelle Auswirkung von Gebäudekandidaten auf das Stadtbild im Kontext von sichtbarkeitsbewusster Stadtplanung und -entwicklung. Wir präsentieren die Visualisierungslösung Vis-A-Ware die es Benutzerinnen erlaubt die Sichtbarkeit von Gebäudekandidaten quantitativ und qualitativ zu evaluieren und zu vergleichen in Bezug auf viele Aussichtspunkte. Vis-A-Ware bietet eine 3D-Ansicht einer urbanen Szene und eine neuartige ‘Ranking’-Ansicht mit der sich Kandidaten und Aussichtspunkte anhand verschiedener

quantifizierter Sichtbarkeitsmetriken vergleichen und filtern lassen. Diese Ansicht ist eng integriert mit anderen Ansichten zur qualitativen Evaluierung und um räumliche Verhältnisse im Stadtbild einschätzen und sich letztendlich für ein kleines Set an Kandidaten entscheiden zu können.

Im dritten Kapitel befassen wir uns mit dem Bereich der geologischen Analyse von digitalen Outcrop (z.Dt. Felsausbiss) Modellen (DOMs) welche eine essentielle Rolle in den derzeitigen Missionen von NASA und ESA spielen, auf der Suche nach Beweisen für die vergangene Existenz von Leben auf dem Mars. Geologinnen interpretieren und vermessen DOMs, erzeugen sedimentäre Protokolle, und kombinieren sie zu Korrelationstabellen. Zur Zeit werden Korrelationstabellen manuell erzeugt. Das ist zeitaufwändig und unflexibel. Mit InCorr präsentieren wir eine Visualisierungslösung die ein 3D-Protokollierungswerkzeug und eine interaktive, datengetriebene Korrelationstabelle umfasst die mit der stratigraphischen Analyse wächst. Korrelationstabellen sind ein wichtiger Bestandteil von geologischen Publikationen und mit InCorr können wir diese Art von Illustration nachvollziehbar machen und den Arbeitsaufwand dahinter drastisch verringern.

Die Ergebnisse dieser Dissertation demonstrieren, dass eine enge Integration von 3D Geometrie- und Attribut-Ansichten essentiell ist für bestimmte Fachbereiche und dass man an diese Integration methodologisch und mit durchdachtem Visualisierungs- und Interaktionsdesign herangehen muss.

# Abstract

In many domains, such as urban planning, civil engineering, or disaster management, analysts need to deal with complex geometric data that also contain multivariate attributes. In addition to the visual analysis of the attribute data, typical tasks involve the localization and understanding of shapes, and judging spatial relations between geometric objects and the surrounding geometry, as for instance a digital terrain model. One way to address this in a visualization design is with coordinated multiple views, combining a 3D geometric view and attribute views by brushing & linking. However, a naive coordination of such views highlights challenges inherent to 3D visualization, as brushed objects may be occluded or lie outside of the current viewing volume. This can easily lead to disorientation and failing of localization, shape understanding, and spatial relation tasks, which ultimately breaks the iterative analysis loop provided through coordinated multiple views.

In this thesis we explore different visual integration approaches for combining geometric and attribute views with respect to three application domains. In the first chapter, we deal with the domain of tunnel inspection and documentation, concerned with revealing patterns in tunnel crack data. We integrate a 3D geometric view with multiple attribute views to a coordinated multiple view solution and present several domain-specific visualization and interaction strategies to overcome the aforementioned challenges. We conclude the chapter with a methodological framework that provides visualization designers with integration guidelines regarding ‘Guided Navigation’, ‘Enhanced Geometric Rendering’, and ‘Similarity-based Analysis’.

In the second chapter, we explore the potential visual impact of candidate buildings to a cityscape in the context of visibility-aware urban planning. We present the visualization system Vis-A-Ware to qualitatively and quantitatively evaluate and compare visibility data of candidate buildings with respect to a large number of viewpoints. Vis-A-Ware features a 3D view of an urban scene and a novel ranking view to compare and filter candidates with respect to visual impact data derived from visibility evaluations. The ranking view is tightly integrated with the other views for qualitative evaluation and to judge spatial relations in the cityscape. We provide users with a workflow to ultimately arrive at a small set of candidates supporting a jury-based decision-making process.

In the third chapter, we are concerned with the domain of geological analysis of digital outcrop models (DOMs) which plays an essential role in the current NASA and ESA missions seeking signs of past life on Mars. Geologists interpret and measure DOMs, create sedimentary logs, and combine them in ‘correlation panels’. Currently, the creation of correlation panels is manual and

therefore time-consuming, and inflexible. With InCorr we present a visualization solution that encompasses a 3D logging tool and an interactive data-driven correlation panel that evolves with the stratigraphic analysis. Correlation panels are an important part of geological publications. With InCorr we provide geologists with an interactive correlation panel that is reproducible and takes significantly less effort to create.

The results of this thesis demonstrate that the tight integration of 3D geometric and attribute views is essential for certain domains and needs to be approached in a methodological way with thoughtful visualization and interaction design.

# Contents

<b>Kurzfassung</b>	<b>ix</b>
<b>Abstract</b>	<b>xi</b>
<b>Contents</b>	<b>xiii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Scope of this Thesis . . . . .	2
1.2 Related Work . . . . .	13
1.3 Contributions of this Thesis . . . . .	17
<b>2 VISAR: A Methodology for Integrating Geometric and Attribute Views</b>	<b>21</b>
2.1 Tunnel Maintenance and Tunnel Crack Analysis . . . . .	22
2.2 Task Analysis . . . . .	23
2.3 On Combining Geometric and Attribute Views . . . . .	24
2.4 Design Goals . . . . .	25
2.5 Visualization and Interaction Design . . . . .	27
2.6 The VISAR Framework . . . . .	34
2.7 Summary . . . . .	37
<b>3 Vis-A-Ware: Integrating Geometric and Attribute Views for Visibility-Aware Urban Planning</b>	<b>39</b>
3.1 Task Analysis . . . . .	40
3.2 Design Goals . . . . .	40
3.3 Related Work . . . . .	41
3.4 Domain Abstraction . . . . .	42
3.5 System Overview . . . . .	43
3.6 Visual Impact Metric Computation . . . . .	44
3.7 Use Case Scenario . . . . .	47
3.8 Visualization and Interaction Design . . . . .	51
3.9 Visibility Analysis . . . . .	58
3.10 Summary . . . . .	63
	xiii

<b>4</b>	<b>InCorr: Interactive Data-Driven Correlation Panels for Digital Outcrop Analysis</b>	<b>67</b>
4.1	Remote Stratigraphic Analysis . . . . .	68
4.2	Design Goals . . . . .	70
4.3	Related Work . . . . .	70
4.4	Geological Analysis of Digital Outcrop Models . . . . .	71
4.5	Design Process and Domain Abstraction . . . . .	76
4.6	Visualization and Interaction Design . . . . .	79
4.7	Evaluation . . . . .	86
4.8	Summary . . . . .	90
<b>5</b>	<b>Conclusion</b>	<b>93</b>
5.1	Summary . . . . .	93
5.2	Outlook . . . . .	94
	<b>List of Figures</b>	<b>97</b>
	<b>List of Tables</b>	<b>103</b>
	<b>Bibliography</b>	<b>105</b>
	<b>Curriculum Vitae</b>	<b>115</b>

# Introduction

In domains driven by 3D spatial data, such as civil engineering, urban planning, or geology, domain experts need to deal with complex geometric data while having to answer intricate spatial questions. For instance, in the case of tunnel maintenance, the structural integrity of a tunnel's concrete shell depends on orientation, length, width, distribution, and moisture of tunnel cracks in the concrete surface. Such geometric data contains multivariate attributes, thus typical tasks require tunnel inspectors to perform traditional multivariate analysis, but also to take into account the location, shape, and extents of 3D objects, and relate them to each other spatially. Analysts have complex geometric questions. For example, urban planners investigate which candidates for a building lot block the view to a landmark the least, or which candidate casts the most shadow onto a neighboring park. Finally, the work of domain experts may involve highly detailed geometric representations acting as essential context to their work. For instance, in stratigraphic analysis, geologists visually compare and annotate different sediment layers in a highly-detailed, textured 3D model of a rock face. This detailed spatial context and multivariate analysis of 3D annotation data are essential to create a regional geological model. Performing such tasks is not supported by a single of-the-shelf tool and often results in domain experts constructing cumbersome workflows where multiple tools are combined through the import and export of files.

This thesis discusses data exhibiting multivariate and complex geometric aspects and its role in an interactive visual analysis process. We explore three different application domains where the focus lies on three different aspects of working with such data. In Chapter 2, we look at integration strategies for geometric and attribute views within a coordinated multiple views setting in the domain of tunnel documentation. In Chapter 3, we use information directly derived from the spatial domain, in the context of 3D visibility analysis in an urban planning scenario. Finally, in Chapter 4, we explore stratigraphic analysis in the domain of sedimentology where annotation and interpretation data originates from a detailed spatial context that needs to be traceable.




Figure 1.1: (a) 3D geometric view showing tunnel cracks as 3D polylines on high-resolution textured tunnel surface reconstruction. (b) Detail view on selected tunnel crack in red. (c) Attribute data of tunnel cracks in the form of rows in a table, with the selected crack highlighted in red. Adapted from [OSP<sup>+</sup>16]

## 1.1 Scope of this Thesis

To illustrate the scope of this thesis, we introduce a scenario from the domain of tunnel documentation. It shall serve as a **prime example** for the nature and complexity of the data and the tasks we want to address. A tunnel surface, which can be tens of kilometers long, is represented by a 3D reconstruction derived from image sequences and laser scans, high-resolution both in geometric and texture detail. Semi-automatic annotation of the image data yields cracks residing on the tunnel surface, described by 3D polylines. In addition to this geometric representation, a crack has a series of attribute values, such as moisture, width, severity category, and orientation with respect to the tunnel's direction (see Figure 1.1). Everything in this 3D scene is georeferenced, existing in a common geospatial context.

A typical task for a tunnel inspector is to judge the structural integrity of a tunnel and advise rehabilitation campaigns accordingly. Therefore, inspectors need to consider geometric and multivariate properties of a large number of cracks and **spatially relate** these properties with respect to the tunnel surface. Figure 1.1 illustrates the geometric and multivariate aspects of cracks within their geospatial context. The image quality of the reconstructed tunnel surface is typically a downgrade from the originally captured images, due to projection and merging






Figure 1.2: (a) Textured 3D reconstruction derived from tunnel surface acquisition by laser scanners and high-resolution cameras. (c) Camera frustum of the acquisition process that contains 3D query point (red) allows users to retrieve (b) the respective unprocessed source image. *The data has been provided by Dibit Messtechnik GmbH.*

of images. Therefore, inspectors, when examining cracks in 3D, also want to inspect the corresponding original images. This requires them to perform **spatial queries** to find out which images contributed to a region of interest based on camera positions and orientations known from the acquisition process, as shown in Figure 1.2. When it comes to supervising a tunnel's excavation, laser scan and image data is acquired from the excavated tunnel heading instead of the tunnel wall. Capturing these so-called tunnel faces in regular intervals leads to a series of 3D rock faces, as visible in Figure 1.3. Engineering geologists analyze, measure, and annotate these 3D tunnel faces to characterize and predict rock properties that will be encountered during excavation. Thus, they **create geometric and multivariate data** from detailed geometric representations and analyze the resulting measurements. Accomplishing such tasks with off-the-shelf tools is either not feasible or results in cumbersome workarounds. On the other hand, creating customized visualization solutions providing users with methods for the visual analysis of geometric and multivariate aspects is challenging.

Throughout this thesis we will present three design studies showcasing a tight integration between geometric and attribute visualizations. To illustrate the relevance of this topic and what we mean by tight integration, we first need to establish the problem space in more detail and examine existing approaches for their suitability to address potential challenges. The remainder of this section is structured into *What?*, *Why?*, and *How?*, first, discussing geometric data faceted with




Figure 1.3: Series of so-called tunnel faces as captured during the advance of the tunnel head. Geological annotations indicate contacts between rock layers. *The data has been provided by Dibat Messtechnik GmbH and has been annotated by Robert Barnes.*

multivariate attributes (Section 1.1.1), secondly, elaborating on the tasks typical for the visual analysis of such data (Section 1.1.2), and lastly, investigating different methods for visualizing such data to accomplish the tasks (Section 1.1.3).

### 1.1.1 What: Multivariate Geometric Data

Munzner [Mun14] defines *geometry* as a dataset type where each item is described by some spatial information in 2D or 3D about its position or shape, as for instance, points, lines, curves, surfaces, regions, or volumes. In this definition items of a geometry dataset do not necessarily have attributes, and therefore attributes are not considered in-depth when discussing the visualization of geometric data. However, as illustrated by our example, attributes are essential for answering questions in certain domains that involve geometric data. We define a data item  $i$  of a geometry dataset with attributes as  $i = (g, a)$ , where  $g$  is a geometric spatial representation and  $a$  is an attribute vector of length  $n$  in the form of  $a = (a_1, \dots, a_n)$ . While 3D points are sufficient to judge the spatial distribution of certain phenomena, more complex spatial questions such as ‘Are two cracks parallel to each other?’ or ‘Do they intersect?’, require an expression of shape beyond a single position. We define  $g_{\text{complex}}$  as a vector of at least two three-dimensional positions in the




Figure 1.4: Grey - Dataset types composed of core datatypes as presented in Munzner [Mun14]. Yellow - ‘Multivariate Geometry’ as new dataset type, differing from ‘Geometry’ by ‘Shape’ and ‘Attributes’.

form of

$$g_{\text{complex}} = (v_1, \dots, v_m) \text{ where } v \in R^3 \text{ and } m \geq 2$$

The spatial representation of an item goes beyond two dimensions and beyond single points. This leads us to a specific dataset type where each item is in the form of

$$i' = (g_{\text{complex}}, a)$$

which we denote as *multivariate geometry*, or *multivariate geometric data*.

Tory and Möller [TM04] identify *spatial data* as data that have a given spatialization, different from data that require a chosen spatialization when visualized, as for instance, locations of cities in contrast to stock prices. In the case of multivariate geometric data, an aspect of each item has an inherent spatial quality, i.e., given spatialization, conveying its location, orientation, and shape. As illustrated in Figure 1.4, another dataset type with a given spatialization are fields, which consist of cells in a grid which have positions and attributes associated to them. This data typically comes from measurements or simulations of continuous phenomena, such as medical imaging or simulations in the field of computational fluid dynamics. In comparison to fields, multivariate geometric data is built on the notion of individual items that have a certain 3D shape. So conceptually, a multivariate geometric dataset itself is not a field, but the shapes of all items might have been derived from a common grid, making the spatial representations of each item a sub-field of the original one. For instance, when segmenting defects in an industrial computed tomography scan of a metal casting, the resulting bubbles exist on a common spatial grid but each item has its individual geometric shape. In the scope of this thesis we are concerned with multivariate geometric data, which are data in form of items comprising a set of attributes and a spatial representation described by a complex 3D shape.

### 1.1.2 Why: Motivation behind Analyzing Multivariate Geometric Data

Spatial data, in the form of geometry or fields, typically occur in the context of tasks that require shape understanding [Mun14]. Shape in this context can be interpreted in two ways, either

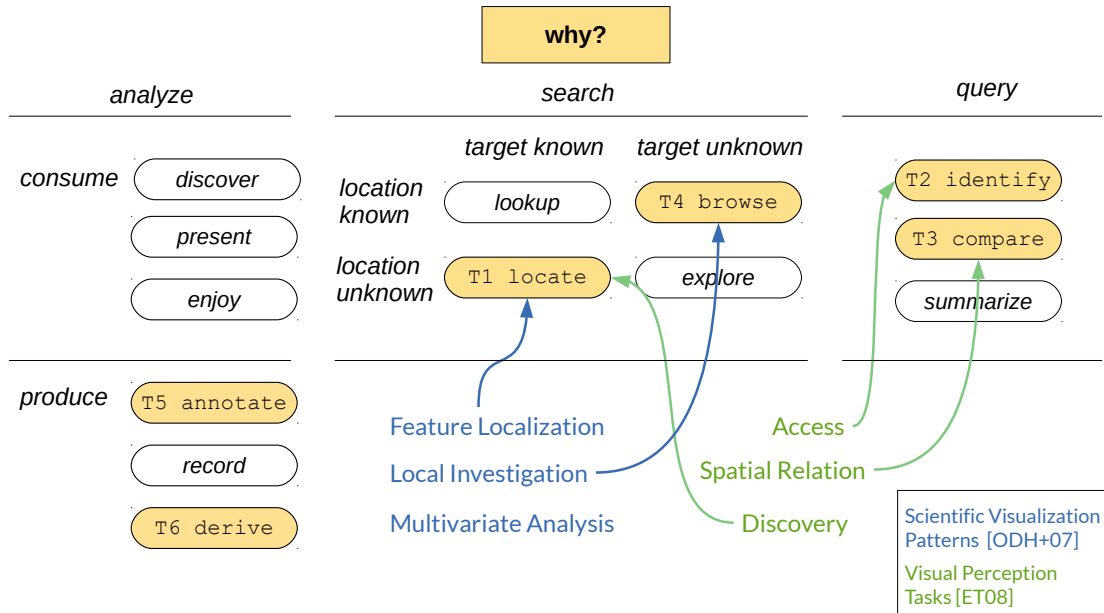


Figure 1.5: Mapping of typical patterns in the visual analysis of scientific data [ODH<sup>+</sup>07] and visual perception tasks [ET08] onto Munzner’s [Mun14] ‘Why: Task Abstraction’ to define tasks for consuming multivariate geometric data (T1, T2, T3, and T4). New multivariate geometric data can be produced either through annotating or deriving (T5 and T6).

(1) individual positions in a dataset make up a certain shape that requires understanding, or (2) each item in the dataset does have a shape that needs to be retrieved, understood, and put into context with other spatial data. Kehrer and Hauser [KH13] show that the visualization methods for scientific data successfully deal with (1) when visualizing multivariate scientific data and provide well-established methods. Based on their work with perfusion data, Oeltze et al. [ODH<sup>+</sup>07] identified the following three typical patterns in visual analysis of multivariate scientific data, which we denote as **scientific visualization patterns**:

- **Feature localization**, the search for spatial data where certain attribute characteristics are present.
- **Multi-variate analysis**, investigation of multi-variate data properties.
- **Local investigation**, inspecting the values of certain attributes with respect to subsets in the spatial domain.

Perfusion data are essentially a 4D spatial field, where for each time step data points consist of an attribute vector and a 3D position. Although this is complex data, the complexity of the spatial component falls below the criterion we defined for multivariate geometric data.

Elmqvist and Tsigas [ET08] provide a taxonomy for 3D occlusion handling techniques. They define three **visual perception tasks** that can be impeded when occlusion occurs. We employ these tasks to express domain tasks that rely on (2):

- **Discovery**, the localization of a geometric representation.
- **Access**, retrieving the shape, color, or any other property beyond the localization of a geometric representation.
- **Spatial relation**, assessing spatial properties of geometric representations in interrelation with other geometric representations.

In summary, the scientific visualization patterns focus on multivariate spatial data which do not include complex spatial shapes, whereas the visual perception tasks consider complex shapes but are not concerned with multivariate data. Multivariate geometric data is practically a union of these two aspects, so we combine the scientific visualization patterns and the visual perception tasks by mapping them onto a common ground, the actions of the ‘Why: Task Abstraction’ presented in Munzner [Mun14]. This generalized mapping, as illustrated by Figure 1.5, allows us to unify the two taxonomies and adapt them to characterize tasks consuming existing multivariate geometric data items.

- T1 *Locate* geometric representations that exhibit certain attribute characteristics matching ‘Feature Localization’ and ‘Discovery’ tasks.
- T2 *Identify* detailed characteristics, such as shape, color, or image texture of a geometric representation which equals to ‘Access’.
- T3 *Compare* geometric representations and their attribute characteristics in 3D space. *Compare* relates to ‘Spatial Relation’ in terms of geometric representations, but also encompasses multivariate aspects.
- T4 *Browse* 3D space in search for geometric representation, matching visual queries, or attribute characteristics. This task is based on ‘Local Investigation’, but also takes geometric characteristics into account.

As producing new data items or new aspects of existing items in a geometric context is not formally addressed by any taxonomy known by the author, we define two additional tasks.

- T5 *Annotate* geometric representations, either in the form of items or context geometry, to create new multivariate or multivariate geometric data. Annotation is mostly performed manually.
- T6 *Derive* new multivariate geometric data through transformations or simulations in the context of 3D geometric shape or geometric questions. Derivation is mostly performed automatically.

The distinction between `T5 annotate` and `T6 derive` can be blurry. For instance, when manually annotating a delineation between rock layers in a geological analysis setting, users pick points on a 3D model. At the same time, a plane fitting algorithm processing the picked points determines the orientation of the underlying geological feature. Thus, the creation of the 3D contact line is manual, through `annotation`, while an orientation attribute is `derived` automatically.

With `T1` to `T6` we propose a *multivariate geometric tasks* taxonomy. This taxonomy particularly addresses tasks, problems, and questions that either involve both geometric and attribute aspects or exclusively concern geometric aspects of the data at hand or when creating new data. ‘Multivariate analysis’ of attribute aspects, as part of the scientific visualization patterns, is still essential in the domains dealing with multivariate geometric data, but the design space of solutions is already discussed in-depth elsewhere [Mun14, Kei02]. Further, detailed characteristics of geometric representations such as shape and texture should play an important role in accomplishing domain tasks, otherwise there is no need to go beyond well-established solutions such as 2D cartographic visualizations or 3D scatter plots and other existing methods presented in Kehrer and Hauser [KH13]. Thus, in the scope of this thesis we are concerned with domain problems that demand a joint analysis of 3D geometric and multivariate aspects and require users to be able to accomplish tasks at the `consuming` level `T1` to `T4` on a regular basis.

### 1.1.3 How: Methods for Visualizing Multivariate Geometric Data

Wang-Baldonado et al. [WBWK00] suggest to use multiple views if there is a diversity in the data, and consequently a single view might result in a cognitive overload. Further, the authors advise to use multiple views if they bring out correlations or disparities in the data. Both guidelines fit well to a combination of spatial and multivariate data. Scientific visualization often combines multiple views to show spatial and multivariate aspects of the data by juxtaposition, and coordinates them via brushing & linking. A selection and highlighting in one view is reflected in highlighting the same data in all other linked views. As stated by Kehrer and Hauser [KH13], coordinating spatial and attribute views in this way allows analysts to perform interactive visual analysis and discover phenomena that may not be apparent in a single view visualization. Munzner [Mun14] summarizes these approaches as ‘facet into multiple views’ with a design space of juxtaposing or superimposing views, sharing encoding, which includes brushing & linking, sharing data, and sharing navigation between views.

Figure 1.6 and Figure 1.7 show an example of a perfusion data analysis by Oeltze et al. [ODH<sup>+</sup>07], where a 3D view is coordinated with scatter plots depicting different multivariate aspects of the data. Users can logically combine selection brushes in attribute views to bring out the desired data points in the spatial view for closer investigation, following the ‘Feature Localization’ pattern. In addition to juxtaposition, Javed and Elmqvist [JE12] define superimposition, overloading, and nesting as possible designs to compose different visualizations. In one of our previous works [SOP<sup>+</sup>15] we provide a taxonomy for categorizing integration techniques to combine spatial and attribute views. We show that most solutions combine these views via sharing selection or data states. That means there is a coordinated highlighting in the form of brushing & linking

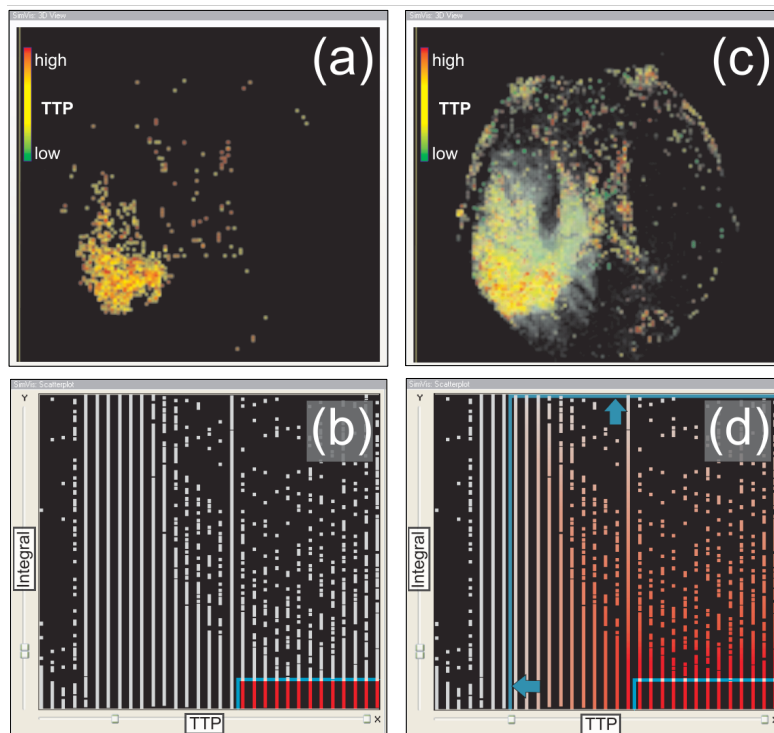


Figure 1.6: Brushing attribute values in scatter plot (b). Revealing ‘infarction core’ in brain in (a) which is color coded by attribute value. Smooth brushing in (d) reveals wider tissue with infarction risk. From [ODH<sup>+</sup>07], simplified caption.

and also data manipulations are shared, as for instance filtering. Further, the authors identify three prevalent coordination patterns:

- **Explore & Feedback**, using interaction in attribute visualizations to explore and understand data presented in spatial visualizations, typical for the visualization of scientific data.
- **Derive & Analyze**, where spatial visualizations are used to generate or query data with respect to spatial properties and then non-spatial views are employed to analyze the resulting data. This pattern is typical for the field of geospatial or cartographic visualization.
- **Balanced Integration**, where neither spatial nor attribute visualizations dominate accomplishing a task, but rather a feedback loop between both types of views is employed.

Recalling the example tasks from the tunnel use case and looking at multivariate geometric tasks suggests that an effective visualization solution requires a balanced integration of geometric and attribute views. To establish a feedback loop or a continuous analysis cycle between geometric and attribute aspects, switching between both types of views should require minimal cognitive effort. Further, it should not be prone to interferences, as for instance, the failing of one of




Figure 1.7: Brain is rotated in (e) for a better overview on infarction in 3D. Smooth brushing in (f) used to convey the shape of the brain. Smooth brushing of other attributes (h) yields similar result in (g) compared to Figure 1.6c. From [ODH<sup>+</sup>07], simplified caption.

the visual perception tasks, loss of context, or spatial disorientation in the geometric view. We argue that ‘shared encoding’ and ‘shared data’ are fundamental to the integration of geometric and attribute views, however, they are not sufficient to reliably provide users with a continuous analysis cycle.

To illustrate the shortcomings of an integration solely based on shared encoding and shared data, we anticipate a scenario that we will revisit in Chapter 2. Consider the following approach for the tunnel use case. As depicted in Figure 1.8, we juxtapose a scatter plot showing crack length and orientation with a 3D geometric view that displays crack polylines and the tunnel surface. Both views share the same data items and the same encoding via brushing & linking. The task at hand requires us to locate cracks on the 3D tunnel surface that have a length of 8 meters or more, and run parallel to the tunnel direction. As illustrated in Figure 1.8a, we brush items in the scatter plot matching our query, and the brushed items are also highlighted in the geometric view (Figure 1.8b). However, brushed items may lie completely or partially outside of the viewing volume of the geometric view (Figure 1.8c). In addition, brushed items may be partially or fully occluded by the geometric representations of other items or other geometry (Figure 1.8d), in this case the tunnel surface. The fundamental tasks of T1 *locate* and T2 *identify* fail fully or partially, requiring users to manually navigate through the 3D scene,






Figure 1.8: (a) 2D scatter plot showing orientation vs. length: linked selection of cracks, which are oriented along the tunnel direction, i.e.,  $0^\circ$ , and which are longer than 8 m. (b) The scene rendered from the current camera position showing only some of the selected cracks (red) (c) Some cracks are partially outside, others are completely outside the view frustum. (d) Cracks occluded by the first tunnel wall are not visible from the current camera position. Adapted from Ortner et al. [OSP<sup>+</sup>16].

which impedes the visual analysis process. To combine geometric and attribute views effectively we propose a *tight integration*, which goes beyond the sharing of data and encoding, as opposed to *standard integration* which solely relies on the sharing of these states. In addition, a tight integration strives to prevent the multivariate geometric tasks from failing to provide users with a continuous analysis cycle.

While attribute views can also be responsible for impeding the analysis process, tasks are more likely to fail in the geometric view as it suffers from occlusion and typically shows only parts of the 3D scene. Further, manual intervention in geometric views, such as 3D navigation or picking, is challenging and prone to disorientation, which often leads to the frustration of the user. In the scope of this thesis we discuss integration approaches which go beyond the sharing of states to address challenges and intricacies that are specific to the integration of geometric

and attribute views. In our use cases, ‘shared navigation’ between geometric and attribute views is not possible in a trivial way as it is between attribute views that have a common axis. There are instances where such views share axes, however, coordinated navigation requires approaches more sophisticated than sharing. We consider this topic as out of scope of this thesis but we will revisit it as an interesting research topic in Section 5.2. In Chapter 2 we provide a methodological framework and guidelines for visualization designers leading to a tight integration letting users continuously explore geometric and multivariate aspects of the data at hand.

### 1.1.4 Design Goals

In the previous sections, we characterized a problem space by defining multivariate geometric data and specifying multivariate geometric tasks T1 to T6, which are typical for working with these data. We argue that a visualization solution operating in this problem space needs to tightly integrate geometric and attribute views. In this section, we derive five design goals for potential visualization solutions, in terms of geometric views, attribute views, and their integration in order to provide users with an effective visual analysis experience concerning multivariate geometric data and the respective tasks. The sophistication necessary for meeting certain design goals strongly depends on the concrete data and domain problem. Solutions may vary with volume and detail of spatial representations, the complexity of the multivariate aspect, and the intricacy of tasks.

**G1 Enable Performance and Quality of 3D Geometric Rendering.** Focusing on the geometric aspect of multivariate geometric data, a solution potentially needs to render large-scale, high-detail geometry. The demand for high-performance geometric rendering may stem from, for instance, a large number of item shapes or high-detail context geometry that is essential for domain tasks. For instance, to render the tunnel surface of our prime example we need to employ out-of-core loading strategies [VM02] because a tunnel dataset of several kilometers length does not fit into main memory. Further, we need to employ level-of-detail rendering techniques [LRC<sup>+</sup>03] because portions of the tunnel we see from one view point may not fit into the graphics card’s memory in full detail.

**G2 Offer Multivariate Analysis Capabilities.** A visualization solution has to provide users with means to perform a multivariate analysis on the attribute aspect of geometric multivariate data. Actual solutions depend on the complexity, clarity, and exploratory nature of domain tasks, and on the form of the multivariate data. Analysis capabilities may include integrating a scatter plot and a histogram, building and integrating novel attribute views, or developing communication interfaces with an existing full-fledged analysis platform. This vast design space is out of scope of this work, but Sedlmair et al. [SMM12] provide a methodology on how to decide which means are appropriate, while numerous methods and examples can be found in Munzner [Mun14].

**G3 Accommodate Computation and Quantification of Spatial Phenomena.** Some domain tasks require the analysis and quantification of varyingly complex spatial questions, essentially deriving new data in the sense of T6. ‘Is a building visible from a certain viewpoint?’ requires a geometric computation and the quantification is only binary. Answering ‘Does a building cast a shadow onto another building at a certain time of day?’ requires a shadow computation. Although

shadow mapping is available in most realtime-rendering engines [Fer05] the challenge lies in the quantification [MDL<sup>+</sup>18]. Finally, finding an answer to ‘Do reflections of sunlight on a building’s glass facade interfere with traffic safety?’ makes it necessary to integrate an elaborate lighting simulation including reflectance models. Depending on the complexity of spatial questions, of their computation and quantification, domain tasks and methods will greatly overlap with the field of visual parameter space analysis, discussed in depth by Sedlmair et al. [SHB<sup>+</sup>14].

**G4 Establish a Common Spatial Context for Geometric Representations.** Investigating the spatial relation of geometric shapes with respect to each other and to the context geometry requires the data to reside in a common spatial context. In geospatial domains, such as surveying, urban planning, or civil engineering, every data item is georeferenced allowing users to retrieve the data in the same geographic coordinate system, typically from a geospatial database. Implementations may range from just having the data already in the same coordinate system, over co-registration and transformations, to integrating an existing spatial database or even building a customized spatial database for a certain use case. An additional complexity arises if data is created from annotation or from simulation as described in **G3**. These produced data must also exist in the same context which may even require to build a customized spatial database.

**G5 Achieve a Tight Integration between Geometric and Attribute Aspects.**

We argue that the integration between geometric and attribute views demands special attention and thoughtful design. We discussed tight integration as an integration between geometric and attribute views, that goes beyond the sharing of states to provide users with a continuous analysis process when changing between views, which essentially concerns the implementations behind **G1** and **G2**. Some use cases require users to derive new multivariate geometric data on-the-fly in the form of starting a new simulation run and immediately compare it to existing runs, both in geometric and attribute aspects, requiring a tight integration between multivariate analysis (**G2**), a simulation engine (**G3**), and displaying the new results in a geometric view (**G1**). With achieving this goal we provide users with a continuous analysis experience when fulfilling multivariate geometric tasks on multivariate geometric data.

## 1.2 Related Work

A number of general-purpose-tools exists that typically achieve one of our design goals (see Section 1.1.4) very well, but are lacking capabilities concerning some of the other goals. 3D realtime-rendering engines are capable of rendering large-scale and high-quality geometric data at interactive frame rates (**G1**), but do not offer analytical capabilities. Despite the fact that they often concern spatial and physical phenomena, such as visibility computation or casting shadows, they rarely provide built-in means for the quantification of these phenomena.

A myriad of commercial and research-based tools for multivariate analysis (**G2**) has emerged over the past decades [Sof13, TIB20, The20, Vis, MFGH08, JJJF07, Wea04]. Some are general purpose, some are more geared towards a certain domain or class of problems, but in essence they combine multiple views via brushing & linking and let users explore the data from different points of view. Most common is the combination of views through juxtaposition but occasionally views

are also superimposed, overloaded, or nested [JE12]. Spatial views in these systems are typically limited to cartographic visualizations or depict a confined 3D spatial phenomenon, as for instance the visualization of a simulated ion trajectory in a cubic ion trap as part of *Improvise* [Wea04]. Roberts [Rob07] provides a survey on coordinated multiple view systems at that time.

For dealing with 3D geometric data and geometric questions, Geographic Information Systems (GISs) are widespread in disciplines such as civil engineering, urban planning, or geosciences. These applications form a compromise between detailed geometric rendering (**G1**) and multivariate analysis (**G2**). They typically either lack complex analytical capabilities [Esr, GRA20, QGI] or are restricted to cartographic views [JJJF07]. A GIS is usually connected to a geospatial database that holds all data in a common geographic context and users can extract these data via spatial queries. Geometric questions such as ‘Is there a direct line of sight between point A and B’ can be answered on a database level (**G3**). GISs appear to at least touch all our goals, and they offer a very elaborate approach to provide a common spatial context (**G4**). However, their general-purpose nature makes it difficult to counteract weaknesses often resulting in cumbersome workflows, such as, reducing volume and detail of geometric data, creating workarounds to quantify spatial queries, or exporting measurement data to multivariate or statistical analysis tools. In many cases in 3D geospatial domains GISs allow users to perform their tasks, but some not in an efficient way. Further, the lacking integration makes exploratory visualization tasks tedious.

In the context of coordinated multiple views for data with spatial aspects, many authors integrated attribute visualizations with 3D views in order to gain insight into the spatial data. They do this by investigating multivariate aspects, as characterized by the ‘Feature Localization’ pattern [ODH<sup>+</sup>07] we discussed in Section 1.1.2. Fuchs and Hauser [FH09], and Kehrer and Hauser [KH13] illustrate how widespread this form of coordinated multiple views is in disciplines such as the visualization of medical images or computational fluid dynamics. These solutions mostly operate on multivariate scientific data with spatial aspects in the form of a scalar, vector, or tensor field.

Solutions providing geometric rendering, in the sense of visualizing geometric shapes of individual items, and means for the analysis of multivariate attributes are often tailored to a specific use case. Jianu et al. [JDL09] combine a geometric view that shows 3D fiber tracts in the brain with a 2D spatial embedding, a color space embedding, and a hierarchical clustering of tracts with brushing and linking. Although the original data consist of voxels, and therefore are a 3D spatial field, the fiber tracts do fulfill the criteria for multivariate geometric data. *Neuomap* by Sorger et al. [SBS<sup>+</sup>13] is another example, in the field of neurobiology. Their work combines a 3D view of the brain of a fruit fly, and a novel graph visualization showing various types of connections between neurons in an abstracted way. The original data are volumetric, but the system supports neuro-scientists in analyzing a variety of derived geometric shapes, such as cell bodies, aborizations, projections, overlaps, and neuropils. These elements are maintained in a spatial database that is optimized to perform overlap queries efficiently.

In the realm of GIS and geospatial data, Chang et al. [CWK<sup>+</sup>07] combine a geometric view with attribute views to explore spatial patterns in demographic data. They also developed their own approach for semantic-aware simplification of a 3D city model preserving urban legibility [CBZ<sup>+</sup>08]. Butkiewicz et al. [BWW<sup>+</sup>08] extend the work of Chang et al. [CWK<sup>+</sup>07]

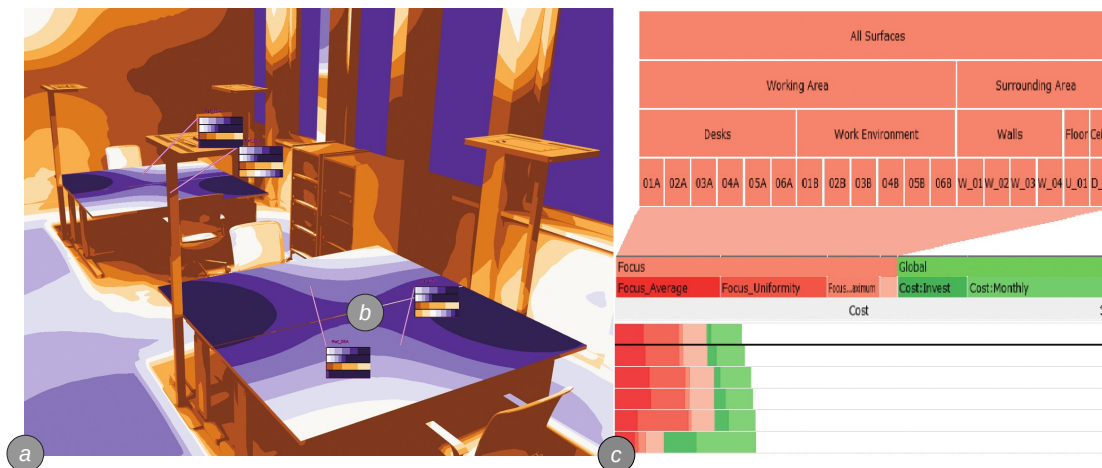


Figure 1.9: Comparing lighting parametrizations according to different metrics in the LiteVis workspace. (a) The geometric view displays a false color rendering encoding illumination quality. For each table, (b) floating annotations show a binned histogram of illumination values from selected simulation runs. (c) The ranking view shows a ranking of simulated lighting parametrizations according to user defined importance values for measurement surfaces and result indicators. From [SOL<sup>+</sup>15], adapted caption.

by a 3D spatial probe metaphor to explore and compare clusters, changing data aggregation via spatial interaction. With Urbane, Ferreira et al. [FLD<sup>+</sup>15] created a framework that allows urban planners to take visibility into account when planning new urban developments. A 3D geometric view displays a cityscape and the potential visual impact of buildings with respect to each other, for instance blocking lines of sight. The visibility simulation is integrated and the impact data can be analyzed in a linked parallel coordinates view. We will discuss an example from the same domain with similar spatial questions in-depth in Chapter 3. Extending the framework of Urbane, Miranda et al. [MDL<sup>+</sup>18] developed Shadow Profiler to allow urban planners to investigate shadows cast in the cityscape working with shadow mapping and raytracing [NM14]. Despite using common computer graphic techniques the quantification of spatial phenomena such as visibility or shadowing requires simulation.

In LiteVis [SOL<sup>+</sup>15] (see Figure 1.9), we combine a geometric view and a ranking view for visual parameter analysis of lighting design solutions. The geometric view (Figure 1.9a) shows a detailed interior scene of an office with measurement surfaces to quantify the light received from an integrated interactive global illumination simulation [LTH<sup>+</sup>13]. Further, stacked histograms (Figure 1.9b) are nested in the geometric view and indicate how well each measurement surface is illuminated. The ranking view (Figure 1.9c) allows lighting designers to see how different simulation runs perform, with respect to a weighted score. Ribicic et al. [RWF<sup>+</sup>13] use visual analysis to evaluate the quality of barrier configurations in simulated flooding scenarios. They combine a large scale geometric view showing terrain, buildings, and infrastructure with scatter plots, histograms, and a parallel coordinates plot. Disaster managers can place barriers interac-

tively in the geometric view and get feedback in terms of how much flooding damage buildings might take. In the work of Cornel et al. [CKS<sup>+</sup>15], in the same scenario, attribute visualizations that show water levels on buildings are directly integrated into the 3D visualization.

In the context of coordinated multiple views, several models and frameworks have emerged. Wang Baldonado et al. [WBWK00] formulate guidelines on when and how to use coordinated multiple views. Boukhelifa et al. [BRR03] present a model on how coordination can be formalized and implemented, while Sorger et al. [SOP<sup>+</sup>15] developed a taxonomy to classify integration techniques of spatial and attribute visualizations. Frameworks include Snap-Together [NS00] and Improvise [Wea04], which are highly focused on the design and coordination of multiple views. Munzner [Mun14] discusses methods and examples on how to facet data into multiple views including shared encoding, shared data, and shared navigation.

We conclude that there are no general-purpose tools that adequately fulfill all the design goals defined in Section 1.1.4 to let users efficiently perform domain tasks as defined in Section 1.1.2. Sorger et al. [SOP<sup>+</sup>15] and Oeltze et al. [ODH<sup>+</sup>07] are two works concerned with the coordination of spatial and attribute views, but not in particular with geometric views displaying multivariate geometric data. Existing approaches and integrations of 3D geometric and attribute views emerged from particular use cases to solve problems in a variety of domains. In this thesis we strive to address and bridge the gap between methodological work and real world use cases for the visual analysis of multivariate geometric data.

## 1.3 Contributions of this Thesis

Sedlmair et al. [SMM12] define "a design study as a project in which visualization researchers analyze a specific real-world problem faced by domain experts, design a visualization system that supports solving this problem, validate the design, and reflect about lessons learned in order to refine visualization design guidelines." The contributions of this thesis are based on three design studies, in which we investigated real-world problems in the domains of tunnel maintenance, urban development, and sedimentology. Together with experts from the respective fields we designed and evaluated solutions and published our results in three research articles. Each of the following three chapters of this thesis is dedicated to a respective design study. In the following we outline the contributions of each article to this thesis. Each research article was a joint effort of multiple authors, whereas the author of this thesis is the first author, which we will refer to in the first person 'I' in the discussion of the individual contributions.

### Contribution 1 - VISAR

In Chapter 2, we present *VISAR: Visual Analytics and Rendering for Tunnel Crack Analysis* [OSP<sup>+</sup>16], a design study for the visual analysis of tunnel cracks in the domain of tunnel monitoring and maintenance. We coordinated a 3D geometric view to depict tunnel cracks and a high-resolution tunnel surface with a scatter plot, a parallel coordinates plot, and a customized aggregation plot for the visual analysis of multivariate aspects of the tunnel crack data. From the resulting coordinated multiple view solution and the nature of the tunnel crack data, we abstracted a general problem space. We address these challenges by domain-specific implementations for the tunnel use case, but derive the methodological framework VISAR (Visual Analysis and Rendering) for handling multivariate geometric data. VISAR builds on top of standard coordinations and introduces the following components specifically addressing the integration of geometric views, 'Localization', 'Overview & Detail', 'Visual Discrimination', 'Occlusion Handling', 'Visual Abstraction', and 'Similarity-based Analysis'. We further provide guidelines for visualization designers leading to a tight integration of geometric and attribute views. In this chapter, we focus on consuming tasks regarding existing items (T1, T2, T3, and T4) and address the design goals **G1**, **G2**, and emphasize solutions to achieve **G5**.

*Thomas Ortner, Johannes Sorger, Harald Piringer, Gerd Hesina, and Eduard Gröller. Visual Analytics and Rendering for Tunnel Crack Analysis. In Proceedings of The Visual Computer, 32(6-8), 859-869. June 2016.*

*Individual Contributions.* I was responsible for compiling the problem space and challenges of combining geometric and attribute views and devised the methodological framework VISAR with its components and guidelines. I developed the use case with a long-term project partner from the field of tunnel maintenance and implemented all domain-specific methods within the framework and conducted design and evaluation sessions with domain experts. I wrote the full article, with regular input from *Johannes Sorger* and *Eduard Gröller*. Furthermore, *Eduard Gröller* was a great help in forming the paper story and reviewing the manuscript. *Johannes Sorger* was involved in continuously shaping the problem space and integration strategies, also complementing our work on the integration taxonomy [SOP<sup>+</sup>15]. He also implemented the communication endpoint for the visual analysis platform. *Harald Piringer* and *Gerd Hesina* had the original idea of combining

3D and 2D visualizations and that there are guidelines missing in doing so. *Harald Piringer* also designed the aggregation plot to fit the tunnel use case and introduced similarity-based analysis to the VISAR framework based on his works in parameter space analysis [BPF11].

### Contribution 2 - Vis-A-Ware

In Chapter 3, we discuss *Vis-A-Ware: Integrating Spatial and Non-Spatial Visualization for Visibility-Aware Urban Planning* [OSS<sup>+</sup>16], providing a design study to qualitatively and quantitatively evaluate, rank, and compare 3D visibility data of candidate buildings in an urban planning scenario. We devised a method to quantify the visual impact of a building with respect to a large number of view points concerning four criteria, landmark-, openness-, and sky occlusion, and building visibility. We created the transposable ranking view, which allows urban planning experts to efficiently compare quantified visual impact data and to identify critical viewpoints and candidates. The developed interaction workflow of *focus*, *filter*, and *transpose* allows them to quickly reduce the set of candidates and relevant viewpoints for decision-making. In addition, the geometric view depicting the cityscape, enhanced with a glyph overlay, and a filmstrip metaphor enables users to judge visual impact spatially and qualitatively. We validated our designs in collaboration with experts from urban surveying, urban planning, and urban design. In this chapter, we derive new data (T6) by quantifying the spatial relations between geometric representations in terms of visibility, addressing the goals **G1**, **G2**, and, with a special focus, **G3**.

*Thomas Ortner, Johannes Sorger, Harald Steinlechner, Gerd Hesina, Harald Piringer, and Eduard Gröller. Vis-A-Ware: Integrating spatial and non-spatial visualization for visibility-aware urban planning. In Proceedings of IEEE Transactions on Visualization and Computer Graphics, 23(2), 1139-1151. January 2016.*

*Individual Contributions.* Supported by *Gerd Hesina*, I got into contact with the urban surveying department in Vienna and could win them as collaborators. I conducted the task analysis, also talking to additional urban planning experts, and acquired domain knowledge in the field of visibility analysis in urban planning. *Johannes Sorger* accompanied me to a number of meetings and collected valuable feedback. I designed and implemented large parts of the 3D visualization including the glyph overlay, the transposable ranking view, and the filmstrip, and I also derived the underlying data model. I also specified the visual impact metrics based on literature review and expert interviews. *Harald Steinlechner* contributed the transparency visualization for displaying multiple building models and the efficient rendering of terrain data. He was also responsible for the messaging protocol between geometric and attribute views and the integration of the HTML/glyph overlay into the geometric view, and the implementations necessary for visibility quantification. *Harald Piringer* and *Gerd Hesina* had the original idea of visibility as a topic worth exploring, also with respect to Harald Piringer's work in visual parameter space analysis [SHB<sup>+</sup>14]. He also had the idea to interact with stacked bar charts via focus and filter, and to encode the focus and context discrimination into the height of the bars. I wrote most of the manuscript with regular revisions from *Johannes Sorger, Harald Piringer, and Eduard Gröller* at different stages of maturity until the final submission.



### Contribution 3 - InCorr

In Chapter 4, we introduce *InCorr: Interactive Data-Driven Correlation Panels for Digital Outcrop Analysis* [OWN<sup>+</sup>21], a design study that supports planetary geologists in deriving regional geological models which are essential to pick promising drill sites for the rover to discover bio-signatures on Mars. Central to InCorr are the design decisions leading to an efficient 3D logging tool and the transformation of a static geological illustration, the correlation panel, into an interactive data-driven version, the InCorrPanel. The logging tool allows geologists to collect geological measurement data through logs. These are densely encoded information visualization summaries of highly detailed geometric representations of rock faces. These logs can then be interactively arranged and compared in the InCorrPanel to derive so-called correlations building the core of regional geological models. We validated our designs against a manual illustration and further we conducted a user-study with a small number of geologists. In this chapter we produce new multivariate geometric data through manual annotation (T5), including automatic processes to derive multivariate attributes with a focus on **G1**, **G2**, and **G3**.

*Thomas Ortner, Andreas Walch, Rebecca Nowak, Robert Barnes, Thomas Höllt, and Eduard Gröller.* InCorr: Interactive Data-Driven Correlation Panels for Digital Outcrop Analysis. *In Proceedings of IEEE Transactions on Visualization and Computer Graphics, preprint, 2021.*

*Individual Contributions.* Through the decade-long collaboration with planetary geologists, I could identify the necessity and opportunity for this design study. I lead the three-stage design process over several years leading to the visualization solution InCorr. I wrote large parts of the manuscript and either realized or oversaw design iterations and collected feedback from our collaborators and other planetary scientists. I devised the algorithm and data model for representing geological logs and also implemented the 3D logging tool. I also produced the use case validation, and designed and conducted the user evaluation. *Andreas Walch* was responsible for revising and integrating the InCorr prototype from the second design stage into PRo3D, also addressing the design questions of alignment and linked interactions. He also cut, annotated, and dubbed the supplemental video. *Rebecca Nowak* was responsible for conducting the second design phase, in which she gathered domain knowledge in terms of semantics and different correlation panel styles defining the minimum viable set of visual encodings, leading to the first InCorrPanel prototype of which many facets are still present in the final version. She also contributed the semantics and the rock list user interface. *Rob Barnes* is a geologists and long time collaborator contributing the domain background of digital outcrop analysis and correlation analysis. He was involved in all design stages, reviewing design iterations, and giving valuable feedback. He was in charge for the geological soundness of the article and also wrote the domain background section. *Thomas Höllt* introduced a methodological approach to the formalization of tasks which contributed to the visualization and interaction design and the respective manuscript sections. He regularly revised the full article and contributed to the related work section. *Eduard Gröller* provided valuable feedback shaping the story of this article over the years and revised the manuscript multiple times.

The contributions of this thesis do not specifically address **G4**, i.e., establishing a common spatial context for geometric representations. This is due to the fact that in the presented design studies our collaborators already provided the multivariate geometric data in a common geospatial context, typically exported from a geospatial database. In Ortner et al. [OTP<sup>+</sup>19], we provide an example where, similar to VISAR in Chapter 2, we integrated a geometric view and a full-fledged visual analysis tool. In addition we interface with a spatial data base. While challenges in such integrations are mostly of an engineering nature, we do see potential design decisions in defining intuitive query interfaces and generating visual proxies in the sense of a progressive visual analysis to hide latencies. However, we consider these topics as outside of the scope of this thesis.

# VISAR: A Methodology for Integrating Geometric and Attribute Views

**This chapter is based on the following publication:**

*Thomas Ortner, Johannes Sorger, Harald Piringer, Gerd Hesina, and Eduard Gröller. Visual analytics and rendering for tunnel crack analysis. In Proceedings of The Visual Computer, 32(6-8), 859-869. June 2016.*

The detection and documentation of cracks in the concrete surface of a tunnel are essential for assessing its condition. These cracks comprise a 3D polyline and several multivariate attribute values, such as length, width, orientation, and moisture. Tasks of analysts are, for instance, to identify patterns which endanger the structural integrity of the tunnel surface or to assess the density of cracks along the tunnel and to identify critical sections. Accomplishing such tasks and evaluating if a repair project is necessary, typically requires the visual analysis of the geometric shape of items, context geometry, and multivariate attributes simultaneously.

The historical workflow in tunnel maintenance involves an analyst inspecting the tunnel surface on-site. Meanwhile inspections are mostly performed virtually on detailed, digitally reconstructed 3D models of the tunnel surface. Tunnel cracks are traced in high-resolution images by semi-automatic crack-detection algorithms. However, the analysis of multivariate data is still mostly performed via spreadsheets and static plots. Since geometric and multivariate data are evaluated separately, no integrated workflow is supported, resulting in tedious work to relate both aspects of the data.

We start this chapter with a design study on how to provide users with visual analysis methods in the context of tunnel maintenance. We then abstract from the specific problem domain of tunnel crack analysis and identify a general problem space emerging from the combination of geometric and attribute views, including obstacles and recurring design questions. Finally, we present VISAR (**V**isual Analytics and **R**endering), a methodology to support visualization designers




Figure 2.1: (a) Reconstructed mesh of the tunnel surface. Tunnels provide an additional metric coordinate system consisting of the value  $S_v$  running along the tunnels axis and  $B$  running perpendicular to it. Each crack comprises (b) its geometric representation, a 3D polyline, and (c) its multivariate attributes, such as length, width, moisture, and orientation. Red highlighting indicates the same crack across (a),(b), and (c).

in creating tightly integrated solutions that allows users to analyze multivariate geometric data effectively.

## 2.1 Tunnel Maintenance and Tunnel Crack Analysis

A tunnel goes through three phases, planning, construction, and maintenance. An essential part of maintenance is the documentation and evaluation of the development of cracks in its concrete surface. Based on these evaluations, maintenance projects are issued to repair or completely replace parts of the concrete surface. Tunnel surveying companies increasingly perform virtual inspections on reconstructed 3D tunnel models instead of doing on-site inspections. This minimizes downtimes of tunnels and enables a more comprehensible analysis. In our case, the surface data is acquired by a rail-mounted laser scanner platform in very high geometric resolution. At the same time, images are taken by high-resolution cameras. Co-registration of geometry and images results in a 3D textured mesh (Figure 2.1a). Ortner et al. [OPH<sup>+</sup>10] provide more detail on how the data is acquired, processed, and rendered. To extract the cracks, the image data is processed by a semi-automatic crack detection algorithm [PCPKS06].

**Tunnel Cracks.** The center of our design study is the visualization of and the interaction with tunnel cracks. The course of each tunnel crack on the tunnel surface is described by a 3D polyline as geometric representation (Figure 2.1b). Along with the geometry, each crack is associated with an attribute vector, consisting of the attributes *length*, *width*, *moisture*, *orientation* (Figure 2.1c), and the additional coordinates  $S_v$  and  $B$ , which we discuss in the following paragraph. Moisture occurs in three categories, *dry*, *moist*, *wet*. Orientation describes a crack's angle in relation to the direction of the tunnel. Its values range from  $0^\circ$  to  $90^\circ$ , where  $0^\circ$  corresponds to cracks parallel to the tunnel direction. Tunnel cracks fulfill the criteria for *multivariate geometric data* (see Section 1.1.1) and exist on the backdrop of a high-resolution 3D reconstruction of a tunnel surface as *context geometry*.

**The  $S_v/B$  Coordinate System.** Experts dealing with linear infrastructures often use a one-dimensional metric value to describe a position along a highway, railroad track, river, or tunnel. This value is often called stationing. In our case, the coordinate  $S_v$  describes meters along the tunnel. With another coordinate  $B$  running on the surface and being orthogonal to the axis, the surface position of every crack can be determined (Figure 2.1a).

**Tunnel Sections.** Modern tunnel surfaces consist of spray concrete, which is wet concrete shot onto the raw excavation of the tunnel, resulting in a smooth surface. The concrete is applied continuously for a certain section of the tunnel, e.g., 120 m along the  $S_v$  coordinate, which forms a homogeneous surface. Individual cracks can be repaired by local injections of concrete, but larger damage typically requires the replacement of the concrete shell of a whole section. Therefore, it is essential to our users to investigate the number of cracks with respect to the  $S_v$  coordinate and identify critical sections.

## 2.2 Task Analysis

Cooperating with experts from tunnel maintenance and tunnel surveying, we identified the following tasks in the context of tunnel crack analysis.

- Locating cracks with anomalous values regarding moisture, length, and orientation.
- Analysis of moisture distribution in the tunnel.
- Identification of critical sections based on the density of cracks.
- Identifying intersecting cracks with high moisture that endanger the structural integrity of the tunnel surface.
- Investigation if moisture in one tunnel tube is also present in an adjacent tube .




Figure 2.2: (a) 2D scatter plot showing orientation vs. length: linked selection of cracks, which are oriented along the tunnel direction, i.e.,  $0^\circ$ , and which are longer than 8 m. (b) The scene rendered from the current camera position showing only some of the selected cracks (red) (c) Some are partially outside, others are completely outside of the view frustum. (d) Cracks occluded by the first tunnel wall are not visible from the current camera position.

### 2.3 On Combining Geometric and Attribute Views

As discussed in Section 1.1.3, to visualize each aspect of the multivariate geometric data we facet it into multiple juxtaposed views and coordinate them by brushing & linking. We visualize the geometric representations of the cracks and the context geometry in a 3D real-time rendering view, which we refer to as the *geometric view*. The attributes of cracks are visualized in views with chosen spatialization, such as scatter plots, parallel coordinates, and histograms, which we denote as *attribute views*.

However, the naive coordination of geometric and attribute views has numerous recurring challenges. Figure 2.2a shows an approach that allows analysts to brush certain cracks in a scatter

plot. In order to judge the spatial distribution of the brushed cracks, they need to identify the corresponding geometric representations in the geometric view. Some of these cracks may fully or partially lie outside (Figure 2.2c) the view frustum of the geometric view (Figure 2.2b), or they may be fully or partially occluded (Figure 2.2d) by other geometric objects in the scene.

Elmqvist and Tsigas [ET08] define three *visual perception tasks*. We denote brushed cracks as part of the *focus* while the remaining cracks and the tunnel geometry are part of the *context* [Hau06]. Applying the task abstraction to the tunnel maintenance use case, the visual perception tasks can be described as follows:

- **Discovery** concerns locating a crack in the geometric view after it has been brushed in an attribute view.
- **Access** is concerned with retrieving the shape or color of a crack in the geometric view.
- **Spatial relation** concerns the assessment of spatial or geometric properties of brushed cracks with respect to other cracks or the tunnel surface geometry.

These three tasks relate to T1 Discover, T2 Identify, and T3 Compare, respectively, as discussed in Section 1.1.2. Throughout this chapter we continue to use the visual perception tasks, as they explicitly express the spatial nature and challenges of geometric aspects.

## 2.4 Design Goals

The situation depicted in Figure 2.2 illustrates that cracks in focus can easily fail the visual perception tasks, due to occlusion or being positioned outside the view frustum. Our visual analysis tool must integrate geometric and attribute views in a way that users are able to perform their domain tasks, while preventing the visual perception tasks from failing. Based on this conclusion, we derived the following design goals:

- **G1.1** Tight integration of the geometric view and the attribute views to allow users a simultaneous visual analysis of both data facets.
- **G1.2** Support the localization of single and multiple cracks in the geometric view based on attribute criteria.
- **G1.3** Encode attribute values in the geometric view to allow users to judge their spatial distribution.
- **G1.4** Identification of clusters or outliers in the geometric and the attribute views.
- **G1.5** Incorporating exploration and visualization metaphors that users are already familiar with.

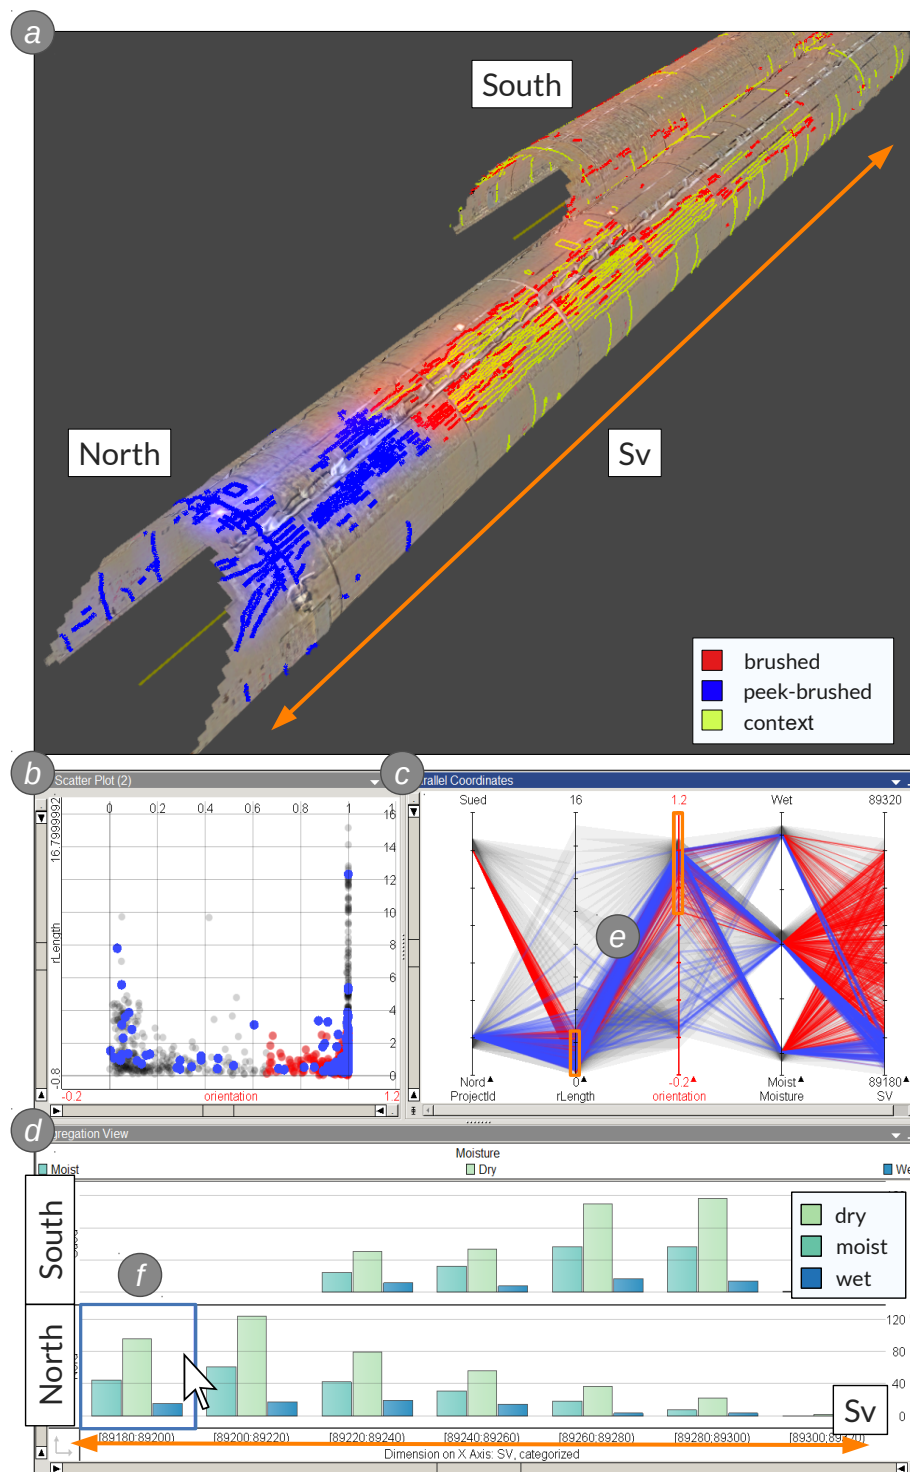


Figure 2.3: System overview: (a) the geometric view shows a textured mesh of the north and south tunnel tube including tunnel cracks as polylines. (b) The scatter plot shows cracks as dots with respect to orientation on the x-axis vs. length on the y-axis. (c) The parallel coordinates plot enables the comparison and identification of trends for many dimensions. (d) The aggregation plot shows the distributions of moisture values in the north and south tunnel tube. Histograms are grouped to 120 m long intervals of Sv. (e) Selections on the length and orientation axis are combined by a logical AND. (f) Peek-brushing over the histogram causes blue highlights on the corresponding cracks in (a), (b), and (c).



## 2.5 Visualization and Interaction Design

In this visualization solution, we combine a geometric view and three attribute views. The geometric view utilizes the 3D real-time rendering framework Aardvark [VRVa], while the attribute views are part of the visual analysis tool Visplore [Vis]. We briefly cover basic coordinations between both parts and discuss the attribute views with a closer look on the aggregation plot. The focus of this design study are the integration-specific design decisions concerning the geometric view. We conclude this section with a ‘Similarity-based Analysis’-tool which appears suitable to explore geometric and attribute data.

### 2.5.1 Linking & Brushing and Color Mapping

Linking & brushing is one of the core concepts of coordinated multiple views and allows users to identify tunnel cracks across views. In our implementation, brushing of cracks in one view results in a red highlighting of the selected cracks, and this highlighting is linked with all other views (Figure 2.3e). Peek-brushing [BP10b] causes a temporary selection of cracks by hovering, which results in a blue highlighting (Figure 2.3f). This allows users to instantly identify and compare tunnel cracks across all views. Users can encode attribute values through color mapping, which is then shared among views. If color encoding is active, we use size and transparency to convey selection and peek-selection. Linking & brushing and shared color encoding are fundamental to achieving **G1.1** and **G1.3**, respectively.

### 2.5.2 Attribute Views

After implementing the aforementioned basic coordinations between the two tools, we discussed potential attribute views with our experts. We agreed on using a scatter plot (Figure 2.3b), a parallel coordinates plot (PCP) (Figure 2.3c), and an aggregation plot (Figure 2.3d). Users found the **scatter plot** very intuitive, and suitable for detecting outliers and clusters with respect to two dimensions. For instance, Figure 2.5a depicts an outlier in a scatter plot of length and orientation, which is a mineral deposit as shown in Figure 2.5c. Users were less familiar with the **PCP**, but welcomed that it offers an instant overview of all relevant data columns. They were especially fond of the possibility to specify arbitrary selection criteria by combining selection intervals on multiple axes, as illustrated in Figure 2.3e. The scatter plot and the PCP partly address **G1.4** since they allow users to identify and brush outliers and clusters.

The **aggregation plot** allows for splitting attribute data hierarchically along two dimensions. The individual parts of the data are presented as a matrix of histograms. Together with our experts, we found this view as suitable for estimating the density of cracks and identifying critical sections. To achieve these tasks, we use the vertical dimension to distinguish between cracks in the south and in the north tunnel tube. We further map the Sv attribute to the horizontal axis and group it into intervals of 120 m, corresponding to the section size of the tunnel. This setup divides cracks into individual tunnel tubes and tunnel sections. After mapping moisture values to the aggregation plot, each section shows individual counts of cracks for each of the three moisture values, as illustrated in Figure 2.3d. Using the Sv coordinate as reference axis addresses **G1.5** and makes the aggregation plot a very intuitive abstraction for domain experts.




Figure 2.4: Visual Discrimination: Different visual variables are manipulated to emphasize the differences between peek-brushed, brushed, and context cracks. (a) Three different widths are used, while cracks in focus have a superimposed glow. (b) Using different widths and colors: blue, red, and yellow-green for peek-brushed, brushed, and other cracks, respectively. (c) Color is used for attribute mapping, while visual discrimination is only conveyed by width and glow.

### 2.5.3 Geometric View

The geometric view provides an interactive rendering of the geometric representations of cracks and the tunnel surface. It allows users to interactively navigate the scene and to assess the spatial extent and distribution of cracks. We visualize the tunnel cracks by using a line shader with screen-space scaling, so each polyline maintains a certain pixel width, regardless of the distance to the viewer. Brushed and peek-brushed cracks (cracks in focus) are highlighted in red and blue (Figure 2.4b), respectively, while the color of context cracks is yellow-green. To further ensure their visibility, brushed and peek-brushed cracks are rendered with a higher pixel width than context cracks. We further use a separated Gaussian blur filter to create a glow effect [Fer04] (Figure 2.4a), which we superimpose onto the cracks in focus. This also preserves visual discrimination of focus and context if color is used to encode attribute values (Figure 2.4c).




Figure 2.5: Localization Transition: (a) Users brush a single crack in the scatter plot. (b) A localization transition is triggered (1) the camera turns to the selected crack, (2) the camera moves to the closest point on a computed orbit sphere (3) the camera rotates around the selected crack and ends up at the characteristic viewpoint inside the tunnel. (c) The scene rendered from the characteristic viewpoint.

Our collaborators tend to use the attribute views to get an overview of the multivariate aspect of their data. They either brush a single crack and seek to inspect its spatial representation (access), or they brush multiple cracks and are interested in their spatial distribution (spatial relation). In both cases, due to occlusion or cracks being outside of the current view frustum, this requires manual 3D navigation which is tedious and can lead to disorientation. To alleviate this and meet design goal **G1.2** we provide *guided navigation* techniques to ensure, that all cracks in focus are inside the view frustum. We use a *virtual x-ray technique* [ET08] and *visual abstraction* [STKD12] to handle occlusion and reduce clutter.

### Guided Navigation

**Single Crack.** Together with our experts we defined what is a characteristic viewpoint [VFSG06] for a single crack, i.e., how should the 3D view provide access to a crack's spatial representation. During a virtual tunnel inspection, analysts investigate individual cracks by 'standing' on the




Figure 2.6: (a) Users investigate the distribution of cracks, which are orthogonal to the tunnel direction. (b) The selection triggers an automatic camera transition to a user-defined overview viewpoint. (c) Clicking on individual cracks seen from the overview viewpoint issues a localization transition, an additional click transitions back to the overview.

tunnel axis (i.e., inside the tunnel surface and 1.70 m above the ground) and viewing them at an almost orthogonal angle, which is based on actual on-site inspection (**G1.5**). Since each crack has an Sv coordinate, we can compute a corresponding position 1.70 m above the tunnel axis. Setting the look-at vector to the center of the crack results in the characteristic viewpoint for a single crack. After users select a single crack in an attribute view (Figure 2.5a), we employ an animated camera transition, called the localization transition, and as illustrated in Figure 2.5b, (1) we animate the camera’s look-at vector to focus on the brushed crack. Before translating, we compute a transitional viewpoint. We create a sphere centered at the crack, where the radius corresponds to the distance between the center of the sphere and the characteristic viewpoint. (2) We compute the closest point on the said sphere as a transitional viewpoint and animate the camera position. (3) The camera is focused on the crack, orbits on the sphere, and reaches the characteristic viewpoint (Figure 2.5c).

**Multiple Cracks.** The aggregation plot provides experts with detailed distributions of attribute values with respect to tunnel sections. However, many tasks require analysts to judge the spatial distribution of attribute values more accurately and to gain immediate access to the corresponding spatial representations. Therefore, we employ camera transitions to allow users to intuitively investigate multiple cracks from overview and detail viewpoints. After users select multiple cracks in an attribute view (Figure 2.6a), which are typically distributed along the tunnel, we trigger an animated camera transition to a user-defined overview viewpoint (Figure 2.6b). Right-clicking on a crack from this position triggers a localization transition. An additional right-click transitions the camera back to the overview. This provides immediate access to detailed spatial representations, for instance, when identifying a cluster of moist cracks.

### Handling Occlusion and Clutter

**Virtual X-Ray.** The localization transition moves the camera to a characteristic viewpoint that is free of occlusion. However, it might be confusing if users are guided to a crack that is occluded during most of the transition. Further, if inspecting the tunnel from an overview position, many cracks are occluded by the tunnel surface. Elmqvist and Tsigas [ET08] present a survey on 3D occlusion management techniques. Based on their categorization, we developed a virtual x-ray technique, which allows users to see cracks in focus through the tunnel geometry. We achieve this by rendering the aforementioned glow for each focus crack without depth testing. Consequently, as it is illustrated in Figure 2.7a, the glows of the cracks in focus shine through the tunnel wall.

**Visual Abstraction.** In some sections of the tunnel, the cracks occur in a high density. Viewing these sections from an overview position, the display is easily cluttered and it becomes difficult for users to distinguish between individual cracks. To counteract this, we replace the geometric shape, the polyline, of a crack with a point sprite if a certain distance threshold is reached. This levels-of-detail approach, or more generally levels-of-abstraction [STKD12], reduces visual clutter and allows users to identify individual crack positions (discovery) and their color (partial access). Consequently, using point sprites as visual abstractions also meets design goals **G1.3** and **G1.4**.

#### 2.5.4 Similarity-based Analysis

In some scenarios, analysts want to compare items with respect to a typical pattern of attribute values. For instance, if there is a dominant pattern of long, dry cracks, oriented along the tunnel direction, analysts are interested in cracks that deviate from this. Therefore we provide a similarity-based analysis tool, which allows users to specify a *focal point*, that is, a point of interest in their data. We quantify similarity by a *distance metric* and treat the resulting distance as a derived attribute value for each crack. Color mapping enables the identification of tunnel cracks that are similar to or deviate from the specified focal point, which serves **G1.4**.

Focal points are typically used in the context of visual parameter-space exploration [BPF11, PBK10]. In general, a focal point is a user-defined  $n$ -tuple specifying concrete values for all or a subset of the  $n$  attributes of a data item. Further, Berger et al. [BPF11] discriminate *global* and *local* updates of a focal point. Local updates only affect a subset of the attributes, while global




Figure 2.7: (a) Brushed cracks are highlighted in red including a glow effect. The glow effect persists even if the respective cracks are occluded. (b) Detailed tunnel crack representations are replaced by point sprites to avoid clutter.

updates affect all attributes at once. For the distance computation, we use a normalized Euclidean distance metric between a crack and the focal point with equally weighted components. Users can select the attributes they want to incorporate into the similarity computation. The focal point can be locally updated by specifying values on the axes of the scatter plot or the axes of the PCP. The respective coordinates of the focal point are represented by green lines. In the geometric view, users can perform a global update by selecting a crack as the focal point resulting in a green highlight.

### 2.5.5 Expert Feedback

We conducted informal feedback sessions for confirming the usefulness of the developed tunnel-crack analysis tool. We interviewed four domain experts: two of whom are from the field of tunnel maintenance (A) and monitoring (B), whereas the other two are from the fields of urban planning (C), and disaster management (D), respectively. Experts A and B were already familiar with using a 3D tunnel visualization for exploration of geometric data. Since multivariate analysis is typically conducted in a paper based, non-interactive form they were eager to specify arbitrary selection criteria in the PCP and the scatter plot and successively refined them.

When investigating multiple cracks, experts A and B found the overview and detail transitions very helpful. All experts considered the localization of focus cracks and the localization transition essential for exploring the geometric view based on attribute criteria. Further, all experts deemed the overview viewpoint and the visual abstractions valuable, since many of their tasks involve the assessment of spatial distributions. Expert C explicitly complimented the implementation of user-specified overview viewpoints. He suggested to add a list for the management of multiple overview viewpoints.

Experts A and B, found the glow implementation for occlusion management helpful in the orientation, but found it confusing during virtual inspection. Consequently, we deactivate the effect if the camera is inside the tunnel. Expert C and D stated that the effect would be useful in an urban scenario, for instance, if objects are hidden behind buildings. For analyzing cracks in a detail viewpoint, two of the four experts desired an orbit navigation mode in addition to the implemented free-fly camera movement.

After a short explanation, experts C and D could see the potential of similarity-based analysis, but they could not immediately imagine how this would translate to their use cases of urban planning or disaster management. The tunnel experts A and B on the other hand could immediately see the benefit of a similarity-based analysis, if applied to the tunnel maintenance scenario. Expert B stated that similarity-based analysis would also translate very well to a tunnel monitoring use case. During construction of a tunnel, segments of it are allowed to shift within a given range of horizontal and vertical movements depending on the type of rock, which surrounds the segment. Considering a large number of shifting measurements along the tunnel over time, it would be helpful to explore them by dissimilarity to normative values. Further, it would be interesting to add time dependent data from deviation measurements and surface geometry of the growing tunnel. In general, our integrated solution was well received. All experts could see the value of a system effectively combining geometric and attribute views on their data. The tunnel experts saw it as a solution to support currently cumbersome tasks.

### 2.5.6 Implementation

We chose to combine two existing systems, the visual analysis tool *Visplore* [Vis] and the 3D real-time rendering engine *Aardvark* [VRVa]. *Visplore* is designed to handle and visualize large amounts of multivariate tabular data. Besides typical 3D rendering features, *Aardvark* can import various data formats and manages out-of-core streaming of large geometry and texture data. Both systems are trimmed to scale well with their respective data facet and were able to

render the scenario at interactive frame rates. Both systems are used separately in numerous application oriented research projects, such as urban planning, tunnel documentation, engine design, and power management. We extended both systems with a custom-built communication layer based on web sockets, which allowed us to handle the coordination between the two systems by exchanging JSON messages.

## 2.6 The VISAR Framework

As outlined in Section 1.2, many authors combine geometric and attribute views through similar integration approaches. The focus of their work, however, typically lies on solving particular domain problems rather than on general applicability. Based on the insights gained during this design study and based on the review of related literature, we could identify common obstacles and recurring design questions inherent to the visualization of multivariate geometric data. As a result, we present VISAR (**V**isual **A**nalytics and **R**endering), a methodological framework addressing this integration on a more general level in order to support visualization designers in effectively combining geometric and attribute views.

Integrated solutions, which are tailored to a specific use case, typically focus on the analysis of certain data items: flooded buildings in disaster management [RWF<sup>+</sup>13], buildings [FLD<sup>+</sup>15] and lines-of-sight [OSS<sup>+</sup>16] in urban planning, illuminated surfaces [SOL<sup>+</sup>15] in interior lighting design, or tunnel cracks in the presented scenario. In Section 1.1.1 we generalized these items as  $i = (g, a)$ , where  $g$  is the geometric representation and  $a = (a_1, \dots, a_n)$  is the attribute vector of length  $n$ . Analogous to the tunnel cracks, items of the form  $i$  are subject to the problem space described in Section 2.3, since  $g$  may lie outside the current view frustum or may be occluded by other geometry in the scene. Although concrete tasks and actual design decisions depend on the specific use case, it is essential that an integrated approach allows the accomplishment of all visual perception tasks.

To support visualization designers in achieving this goal, our methodological framework addresses the integration between heterogeneous systems on two levels, which are reflected in the two-layer-structure of VISAR: the *mirroring layer* and the *integration layer* (Figure 2.8). The mirroring layer covers coordinations between views in general, while the integration layer explicitly deals with the coordination of geometric and attribute views in order to prevent visual perception tasks from failing.

### 2.6.1 Mirroring Layer & Integration Layer

The **mirroring layer** is responsible for coordinations, which are shared among all views regardless of their type. This includes the selection, i.e., brushing, of items or the encoding of attribute values with colors. Since the mirroring layer encompasses principles implemented in most CMV systems, it acts as the foundation of the integration layer. This includes coordinated interactions, such as linking & brushing [FH09], focus & context [KH13, TBPD11], or shared color encoding [JDL09, GRW<sup>+</sup>00].






Figure 2.8: The VISAR framework is divided into two layers: the mirroring layer and the integration layer. The mirroring layer contains simple coordinations, such as ‘Selection’, ‘Peek-Selection’, and ‘Color’. The components of the integration layer are concerned with more complex coordinations that facilitate the visual perception tasks (see Section 2.3): ‘Guided Navigation’, ‘Enhanced Geometric Rendering’, and ‘Similarity-based Analysis’.

The **integration layer** is concerned with the coordination between geometric and attribute views. Its components explicitly address the discussed problem space in order to facilitate the visual perception tasks. Through literature review and generalization of the implementations discussed in Section 2.5, we derived the components ‘*Guided Navigation*’, ‘*Enhanced Geometric Rendering*’, and ‘*Similarity-based Analysis*’ for the integration layer. We will elaborate on each component and its sub-components adhering to the following structure: Purpose of the component, design goals of its sub-components, design choices, and comparison to literature.

### 2.6.2 Guided Navigation

The purpose of the ‘Guided Navigation’ component is to allow users the localization of the geometric representations of items in focus by means of automated camera transitions. The intent of localization is typically preceded by brushing one or more items in an attribute view. We distinguish two goals:

1. Provide users with the discovery and detailed shape access of a single item.

2. Provide users with the discovery of multiple items and allow users to judge their spatial relation.

**Localization.** We address (1) by a localization transition that animates the camera to a characteristic viewpoint. The transition itself and the characteristic viewpoint typically depend on the use case. Design choices lie in the computation of a characteristic viewpoint and the definition of a transition path that does not cause disorientation. Viola et al. [VFSG06] use an information theory approach to estimate characteristic viewpoints of internal organs in a medical data set. When selecting an organ, a localization transition is triggered that moves the camera on a bounding sphere. Buchholz et al. [BBD05] discuss a constraint-based framework for navigation in virtual 3D landscapes.

**Overview & Detail.** We address (2) by an overview transition that animates the camera to reach an overview viewpoint, which contains all items in focus. Design choices concern the definition of a suitable overview viewpoint and how to provide a transition that prevents users from disorientation. Although there is no example actually computing a characteristic overview viewpoint for a changing set of items in focus, approaches can be found in the category of tour planning techniques by Elmqvist and Tsigas [ET08].

### 2.6.3 Enhanced Geometric Rendering

The ‘Enhanced Geometric Rendering’ component manipulates the geometric representations of items in order to achieve the following goals:

1. Discrimination between items in focus and items or other geometry as part of the context.
2. Completing all visual perception tasks for items in focus despite suffering from occlusion.
3. Allowing users to judge spatial relation and provide access to encoded values from overview viewpoints.

**Visual Discrimination.** In general, goal (1) is achieved by highlighting geometry in focus or lowlighting geometry belonging to the context [Hau06]. In many cases, the highlighting or lowlighting of the geometric representation of an item requires design decisions that preserve color, texture, size, or shape, which may be essential to an effective analysis. Trapp et al. [TBPD11] evaluated several techniques for highlighting in 3D geovirtual environments. Outline-based techniques and style-variant techniques, that consider use-case context, are suited best to enable visual discrimination and preserve relevant spatial properties.

**Occlusion Handling.** To address goal (2), a suitable occlusion management technique is necessary to ensure that discovery, access, and spatial relation tasks do not fail, although the items in focus are occluded. Elmqvist and Tsigas [ET08] provide a wide variety of design alternatives. Virtual X-ray techniques appear to be most suitable for geometric views combined with attribute views, since they are reliable and preserve most of the geometric properties.

**Visual Abstraction.** With increasing distance, and also dependent on the density with which items occur, the display becomes cluttered and discovery is impeded. We address goal (3) by replacing the original geometric representations of items by visual abstractions. Such an abstraction may range from a geometric simplification, to a glyph representation summarizing multiple items [OSS<sup>+</sup>16], up to small attribute views integrated into the 3D scenes [SOL<sup>+</sup>15, BWW<sup>+</sup>08]. Chang et al. [CWK<sup>+</sup>07] employ a hierarchical simplification of city models adhering to perception constraints derived from urban planning [CBZ<sup>+</sup>06]. Semmo et al. [STKD12] developed a levels-of-abstraction framework for blending between discrete simplification levels for various geospatial features.

#### 2.6.4 Similarity-based Analysis

The ‘Similarity-based Analysis’ component allows users to explore their data by means of similarity and deviation with the following goals:

1. Comparison of items with respect to a user-specified focal point in the data.
2. Detecting clusters of similar items or identification of items deviating from a common pattern.

**Distance Metric.** We define the similarity of an item  $i$  to the focal point  $f$  as follows:

$$similarity(f, i) = \left( \sum_{j=1}^n w_j \cdot dist(f_j, i_j)^p \right)^{\frac{1}{p}} \quad (2.1)$$

For  $dist(f, i)$  different distance metrics can be inserted, such as the Euclidean distance ( $p = 2$ ). Further, different weights can be applied in form of the  $n$  dimensional weighting vector  $w$ . Which distance metric to use, depends on the application scenario and the involved data types. Migut et al. [MvGW11] present approaches for metrics considering nominal attributes and weights. Mesh comparison metrics, as presented by Schmidt et al. [SPA<sup>+</sup>14], allow visualization designers to integrate geometric properties of the data into the similarity computation.

**Focal Point Updates.** Berger et al. [BPFG11] distinguish *global* and *local* updates of the focal point. An integrated solution may allow users local or global updates of the focal point in geometric or attribute views. Selecting a particular entity as the focal point in any view is a straightforward way for a global update. In Legible Cities [CWK<sup>+</sup>07], users can select individual neighborhoods as focal point in 3D by a pin metaphor in a census dataset.

## 2.7 Summary

We developed an effective solution for the visual analysis of tunnel cracks by integrating geometric and attribute views. We derived VISAR, which we see as an extended discussion of our design study in terms of generalization and believe in its applicability to a wide range of use cases involving multivariate geometric data. At this point we want to briefly reflect on the limitations and possible extensions of this methodological framework.

The VISAR methodology mostly concerns the side of the geometric visualization, which is based on the fact that the three visual perception tasks are more likely to fail in 3D space than in the visualization space of the attribute visualizations. Furthermore, our users tended to employ attribute views for exploring the geometric view, which is denoted as ‘Explore & Feedback’ [SOP<sup>+</sup>15]. However, we realized, that the visual feedback of, for instance, scatter plots reacting to a selection in the geometric view is often not sufficient to attract a users’ visual attention. This issue could be alleviated by using visual links [WPL<sup>+</sup>10] connecting geometric representations and attribute representations. In general, the presented methodology strives to generate solutions that provide users with a balanced integration in order to provide users with a continuous feedback loop between both visualization domains [SOP<sup>+</sup>15].

# Vis-A-Ware: Integrating Geometric and Attribute Views for Visibility-Aware Urban Planning

**This chapter is based on the following publication:**

*Thomas Ortner, Johannes Sorger, Harald Steinlechner, Gerd Hesina, Harald Piringer, and Eduard Gröller. Vis-A-Ware: Integrating spatial and non-spatial visualization for visibility-aware urban planning. In Proceedings of IEEE Transactions on Visualization and Computer Graphics, 23(2), 1139-1151. January 2016.*

The designation of urban areas for construction is typically followed by an open call, where architecture firms submit building designs. In case of a redesignation of inner city areas, it is common that such a call yields up to 30 candidate buildings. A new building in such an area possibly affects the cityscape, and typically sparks controversy between various stakeholders, such as neighbors, the property developer, and decision-makers from politics and urban development.

Although there are many factors to consider when comparing candidates, visibility analysis plays an increasingly important role especially in controversial cases, for instance, if the occlusion of a landmark might occur. Typically, experts from urban surveying conduct the following series of analyses to support an informed discussion between decision-makers: creating visibility and landmark occlusion maps for each candidate, photomontages overlaying real images with virtual candidate buildings, and offline renderings for each building from a few viewpoints. Further, haptic models of each candidate and of the surroundings are created to explore different scenarios during a jury discussion.

In this chapter we discuss the design study of Vis-A-Ware, a tool tailored to the 3D visibility analysis of building alternatives in urban environments. We illustrate the applicability of Vis-A-Ware based on an artificial use-case scenario and present feedback by experts from urban

planning, urban development, and urban surveying. As a secondary contribution we present the quantification of visual impact by combining multiple metrics derived from the urban planning literature [YPL07].

In the current visibility analysis process, the impact of candidates on the cityscape is mostly judged by aesthetics and visual appeal, based on montages, renderings, and the haptic model. Apart from a potential subjective bias, a pure qualitative assessment restricts the analysis process to a small number of viewpoints. For a quantitative analysis, geographic information systems (GIS) are used. For example, in landmark occlusion analysis, a mask is computed that indicates areas in a city from where a landmark is visible. Computing 2D visibility maps for up to 30 candidates is very time-consuming and GIS tools do not offer functionality to compare dozens of map layers efficiently. Furthermore, it is cumbersome to extend the described method to a more granular quantification of visibility beyond 'landmark is visible' and 'landmark is occluded'.

### 3.1 Task Analysis

Cooperating with experts from urban surveying, urban planning, and urban development we identified the following tasks in the context of 3D visibility analysis.

- Line-of-sight and landmark occlusion analyses. Investigating for a number of candidates if they block the view to a landmark.
- Selection of relevant viewpoints. Offline renderings and montages are expensive to create, so a careful selection of viewpoints is necessary.
- Analysis of the visual impact on the cityscape. High-level goal of determining what impact a certain candidate has on the cityscape.
- Comparison of candidates with respect to visual impact. High-level goal of comparing different aspects of candidates in terms of their visual impact.

Relating these domain tasks to the multivariate geometric tasks (Section 1.1.2), the `T6 Derivation`, that is quantification, of visual impact data is essential. It provides users with a quantitative comparison, and also a qualitative comparison in the form of `T2 Identifying geometric features of candidates`. `T3 Comparing candidates with one another and spatially relating them in the cityscape` are essential for accomplishing domain tasks.

### 3.2 Design Goals

Based on the previous observations, we derived the following design goals, for a visualization tool supporting experts in the aforementioned domain tasks.

- **G2.1:** Compute intuitive metrics for quantifying the visual impact of candidates with respect to specific viewpoints.

- **G2.2:** Tight integration of geometric and attribute views to allow for a linked analysis of quantitative and qualitative data.
- **G2.3:** Fast identification of candidates or viewpoints exhibiting high visual impact values.
- **G2.4:** Provide users with an overview of the spatial distribution of viewpoints with high visual impact.
- **G2.5:** Intuitive filtering, ranking, and comparison of candidates as well as viewpoints.
- **G2.6:** Incorporating exploration and visualization metaphors users are already familiar with from existing tools.

### 3.3 Related Work

#### 3.3.1 Visibility Analysis

Determining the visibility of 3D objects is an essential and well-researched topic in the field of real-time rendering. The most important application is occlusion culling in order to avoid drawing invisible objects. This reduces potential overdraw in complex 3D scenes [AMHH11]. As presented by Bittner et al. [BWPP04], most occlusion culling implementations spatially subdivide the 3D scene. Brunnhuber et al. [BSFL<sup>+</sup>12] present a system to evaluate signage for pedestrians inside large buildings. They use projections, similar to shadow mapping, originating from simulated pedestrians to determine how good signs are potentially perceivable by a person. Although they use a fine-grained concept of visibility, their tool does not offer any analysis capabilities beyond false-color mapping. Visibility analyses are a standard feature in many geographic information system (GIS) applications, such as ArcGIS [Esr]. Typical tools include line-of-sight intersections for urban scenes, or viewshed analysis in rural areas. These computations are often time-consuming and are therefore performed offline, while results are encoded into map layers. Geoweb3D [Geo] alleviates this rigidity and offers interactive viewshed evaluations in a GIS environment by using a technique similar to Brunnhuber et al. [BSFL<sup>+</sup>12]. None of these tools allow experts the comparison of multiple viewpoints, or to determine visibility values for individual buildings. Doraiswamy et al. [DFL<sup>+</sup>15] developed several metrics to quantify the quality of floor views in order to create efficient building designs in terms of effort and real-estate value. Their system focuses on the optimization of design parameters to support the iterative design process between architects and developer clients.

#### 3.3.2 Decision-Making & Ranking

Multiple-Criteria Decision Analysis (MCDA) refers to making a decision in the presence of multiple, usually conflicting criteria. Many authors deal with MCDA in the context of multi-objective optimizations [BP10a, TMH09]. The described methods are typically applied to multiple-criteria design problems such as in visual parameter-space exploration, where a large or infinite number of design alternatives is present [RWF<sup>+</sup>13, WFR<sup>+</sup>10, SHB<sup>+</sup>14].

In contrast, multiple-criteria evaluation problems contain a finite number of alternatives, which is often the case in geospatial analysis, such as automatic MCDA analysis in the context of land-suitability analysis [JZ10]. Andrienko and Andrienko [AA03] extend a GIS application with attribute views in order to gain insight into automatic MCDA algorithms. Bruckner and Möller [BM10] cluster visual results from an infinite design space into distinct finite alternatives. Sorger et al. [SOL<sup>+</sup>15] use a weighted ranking view to evaluate lighting setups for MCDA in the context of lighting design. Gratzl et al. [GLG<sup>+</sup>13] present a tool called LineUp to intuitively rank a large number of alternatives. While LineUp focuses on addressing ranking problems in general by weighted scores, our work concentrates on ranking and efficiently exploring distributions in an urban-planning use-case. Our ranking approach and design choices were inspired by their work.

Several authors present tools to support decision-making in the context of urban planning and urban development. Andrienko et al. [AAB<sup>+</sup>10] present a framework to aggregate and analyze large sets of spatio-temporal data regarding traffic and crime analysis. Zhang et al. [ZYM<sup>+</sup>14] present a framework for detecting, aggregating, and analyzing public utility service problems through coordinated multiple views. Zeng et al. [ZFA<sup>+</sup>14] quantified mobility data and developed a novel visualization, letting users explore and compare branches in a public transportation system. Ferreira et al. [FLD<sup>+</sup>15] aim to provide an application that gives a holistic overview of urban data, including visibility calculations. Compared to our approach, their work does not offer a way to efficiently compare multiple building candidates. It also does not consider the special requirements that urban planners have to take into account when approving a new building, as for instance, multiple lines of sights and signature viewpoints.

## 3.4 Domain Abstraction

If a municipal area is designated for a construction project, it is statutory for the local government to initiate a call for proposals. Multiple architecture offices submit their proposed buildings, which we will refer to as *candidates*. For prestigious projects such a call can yield up to 30 candidates. These candidates are then checked against hard criteria, such as maximum height, minimum energy efficiency, minimum floor space, or office to apartment ratio. Since these criteria are known by the architecture offices beforehand, only a small number of candidates is discarded in this process. Besides meeting these hard constraints, a candidate is required to 'visually fit' into the cityscape.

For this purpose, urban planning experts produce a visibility analysis, often consisting of *line-of-sight* analyses, i.e., rendered still images from certain *viewpoints* in the city, and photomontages overlaying real images with candidate renderings. Lines of sight are point to point connections from vantage points to typically a *landmark*, for instance, the line of sight from a hill top to a cathedral. If a candidate severely occludes a landmark, it is likely to be rejected. Additionally, urban planning experts select a handful of relevant viewpoints. Such viewpoints are chosen through experience and are then used for rendering still images and creating photomontages. The resulting images support a qualitative evaluation in a jury discussion.

So far candidates violating hard criteria were already discarded, but other non-visibility parameters of the remaining candidates still factor into a final decision. In this regard, the partitioning of



floor space into offices, apartments, shops, or hotel rooms is relevant since potential revenues can be derived from that. However, we consider these parameters as out-of-scope of Vis-A-Ware for two reasons: First, our collaborators are only concerned with the visibility analysis of candidates and could not provide these data. Second, algorithms and tools for MCDA of parameters not concerning 3D visibility are already widely in use [AA03, JZ10].

Factors that are not directly related to the building design itself are outside of the problem space as well, since the site for development is already chosen. Examples are the connection to public transport [ZFA<sup>+</sup>14], the access to public infrastructure (e.g., shops, schools, museums, etc.), or the proximity to recreational areas [FLD<sup>+</sup>15].

### Quantifying Visual Impact

A qualitative analysis of  $n$  candidates with respect to  $m$  viewpoints involves  $n * m$  comparisons. To support this process with visual analytics methods, we need to quantify the visual impact that candidates cause on specific viewpoints. Urban planning researchers have sought to quantify the visual impact based on the idea that geometric properties relate to the perception of urban space [YPL07]. Following Viewsphere, introduced by Yang et al. [YPL07], we derived four visual impact metrics (VIMs): *landmark occlusion*, *candidate visibility*, *openness occlusion*, and *sky occlusion*. We discussed the VIMs with urban planning experts, who confirmed that these metrics are well-suited for quantifying the visual impact. The described metrics are derived from false color images, as illustrated in Figure 3.2c. In the following we provide a definition of VIMs while their computation will be described in Section 3.6.

- **Landmark Occlusion.** Image area of landmarks occluded by a candidate.
- **Candidate Visibility.** Image area covered by a candidate. Describes how prominent a candidate is from a specific viewpoint.
- **Openness Occlusion.** Image area of open space (with respect to existing buildings and terrain) covered by a candidate. Openness is based on the Spatial Openness Index presented by Yang et al. [YPL07].
- **Sky Occlusion.** Image area of the sky covered by a candidate.

## 3.5 System Overview

After the domain abstraction, we are confronted with a multiple-criteria decision problem where users compare and evaluate a finite set of solutions. Due to the mixture of qualitative (aesthetic appeal and perception) and quantitative aspects (VIM values) in this scenario, there is typically no single optimal candidate. The goal of Vis-A-Ware is to support urban planners in reducing the set of candidates and viewpoints effectively and efficiently without relying on time-consuming offline renderings or GIS analyses. The resulting reduced set is then presented to decision-makers as a foundation for an informed jury discussion.

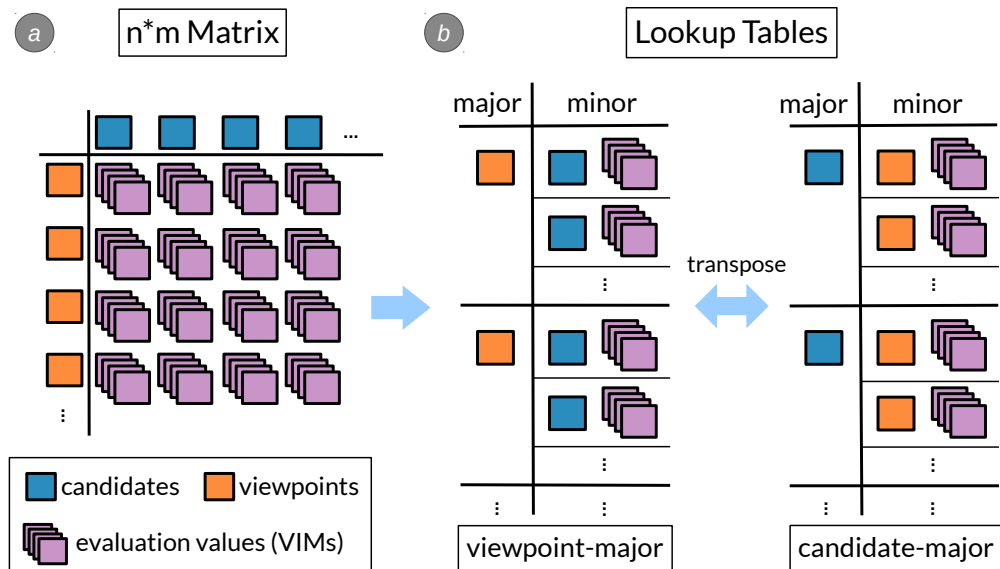


Figure 3.1: Data Model. (a) Evaluations are represented by an  $n \times m$  matrix of candidates and viewpoints. Each combination holds a series of visual impact metric (VIM) values. (b) This representation is transformed into a lookup table either governed by viewpoints (viewpoint-major) or by candidates (candidate-major).

We address the problem of visibility-aware decision-making in several steps. From the domain abstraction we derive a flexible data model which represents candidates, viewpoints, and VIMs (Section 3.6.1). We developed an image-based approach to quantify visibility and quickly evaluate candidates for hundreds of viewpoints (Section 3.6.2). We process the resulting VIM data by normalization and binning to make it more intuitive and comparable (Section 3.6.3). We constructed a use case consisting of 20 candidate buildings and 53 viewpoints (Section 3.7). We provide a novel ranking view (Section 3.8.1) to explore the processed VIM data through means of ranking, focus and context, and detail-on-demand. The ranking view is linked with a 3D view (Section 3.8.2) for assessing spatial distributions and an image viewer (Section 3.8.3) for qualitative comparisons.

## 3.6 Visual Impact Metric Computation

### 3.6.1 Data Model

Our data model consists of the following items: *candidates*, *viewpoints*, and *VIMs*. The combination of a viewpoint, a candidate, and a set of VIMs we denote as an *evaluation*. Each evaluation is associated with a rendered image of the scene. The entire evaluation data can be organized in an  $n \times m$  matrix, where  $n$  and  $m$  are the numbers of candidates and viewpoints, respectively (see Figure 3.1a). In our case, each cell of this matrix contains a vector of four VIM values for landmark occlusion, candidate visibility, openness occlusion, and sky occlusion.

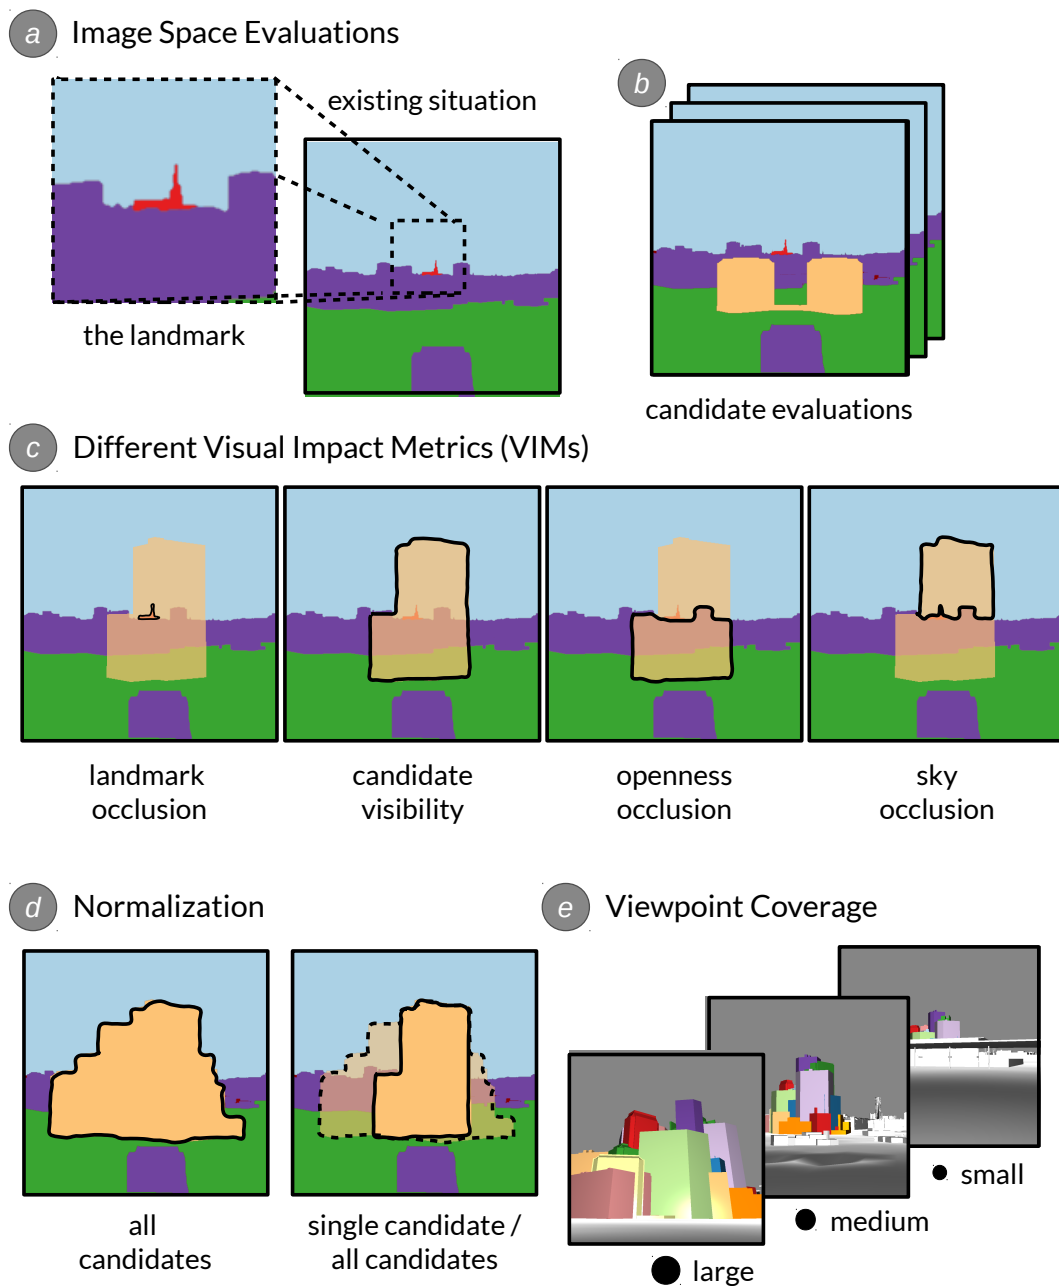


Figure 3.2: False color image evaluation of (a) existing situation, (b) for each candidate. (c) Four metrics to quantify visual impact: landmark occlusion, visibility, openness, sky occlusion. (d) ‘All candidates’ evaluation for normalization of metrics (except for the landmark occlusion metric). (e) Ratio of ‘all candidates’ to image size: large, medium, small.

From this data layout we derive two lookup tables as shown in Figure 3.1b. The table on the left shows candidate scores for each viewpoint (*viewpoint-major*), while the table on the right shows impacted viewpoints for each candidate (*candidate-major*). We denote the items by which a row of the lookup-table is identified as *major*, whereas the items that form the content of each row are denoted as *minor*. Switching, from a viewpoint-major to a candidate-major representation or vice-versa, will be referred to as *transposing*.

For our approach of a quantitative visibility analysis, we interpret viewpoints and lines of sight as measurement probes for visibility in a 3D scene. We define each viewpoint by its spatial position, orientation, and horizontal field of view. We model lines of sight as a special case of a viewpoint with a narrow field of view. The cityscape, where the visibility analysis is performed, is represented through a 3D city model and a digital elevation model of the terrain.

#### 3.6.2 Image-Based Approach

For evaluating the visual impact that candidates cause on a viewpoint, we render the view from the respective viewpoint in false colors, as illustrated in Figure 3.2. We turned off lighting and shading and defined unique colors for each of the following components of the scene: landmarks (red), existing buildings (purple), terrain (green), sky (blue), and candidate buildings (orange). We render a single false-color image for each candidate (Figure 3.2b). We further render an image containing the *existing situation* without any candidate building (Figure 3.2a). This lets us compute the difference of each VIM for each candidate in comparison to the existing situation. Finally, we render an image with all candidates at once (Figure 3.2d), which we use for normalization, as will be described in the next subsection. After each rendering, the pixels belonging to each unique color are counted. As a result, we get the area each component covers in the rendering from the current viewpoint. In addition to the false-color renderings, we save a rendered image from the current viewpoint and candidate combination for the qualitative analysis. The presented approach is straightforward to implement, fast, and offers a pixel-accurate visibility-evaluation method. The current implementation summarizes all landmarks into a single score. Our approach is also compatible with handling multiple distinct landmark scores. This would require an additional VIM per landmark in the data table and an additional color in the false-color rendering.

#### 3.6.3 Normalization and Binning

Our image-based evaluation approach yields numbers of pixels. Pixel counts as raw output values are unintuitive and since different parts of the scene naturally cover vastly different areas (e.g., sky versus existing buildings), it is difficult to compare candidates or viewpoints across metrics. Therefore we decided to normalize our VIM data. This is straightforward for the landmark occlusion metric ( $lom$ ):

$$lom_i = (lom_0 - lom_i) / lom_0 \quad (3.1)$$

$\widehat{\text{lom}}_i$  is the ratio of the landmark area occluded by a candidate  $i$  ( $i = 1 \dots n$ ). It is described by the difference of landmark pixels of the existing situation ( $\text{lom}_0$ ) and landmark pixels in the image containing candidate  $i$  ( $\text{lom}_i$ ) divided by  $\text{lom}_0$ .

Our users confirmed that it is intuitive to measure landmark occlusion in percentage of landmark area occluded by a candidate. However, we cannot normalize openness occlusion and sky occlusion in the same way, since they only yield very small relative differences which are difficult to interpret. Both are also very sensitive to the viewpoint orientation. For instance, if tilting the viewpoint slightly upwards, the number of sky pixels increases substantially, while the number of existing-building- and terrain-pixels decreases. Since there is no candidate building in the existing situation, we tried to measure the visibility of a candidate relative to the total number of image pixels, which as well did not yield intuitive results. To address this we compute the normalized VIM value  $\widehat{\text{vim}}_i$  for candidate  $i$  in a given viewpoint as follows:

$$\widehat{\text{vim}}_i = (\text{vim}_0 - \text{vim}_i) / (\text{vim}_0 - \text{vim}_a), \quad (3.2)$$

where  $\text{vim}_0$  is the number of pixels of a VIM in the existing situation,  $\text{vim}_i$  is the number of pixels in the ‘single candidate’ image containing candidate  $i$ , and  $\text{vim}_a$  is the number of pixels in the ‘all candidates’ image, an image containing the footprint of all candidates together, as illustrated in Figure 3.2d. In case of the candidate visibility metric the value of  $\text{vim}_0$  equals 0. This is due to the fact that there is no candidate building in the existing situation. For computing the openness occlusion metric, we add the pixel counts of the existing buildings and the terrain in the respective images.

After normalization, the VIM values are scaled to percentages and are binned into four classes:  $[0 - 25[$ ,  $[25 - 50[$ ,  $[50 - 75[$  and  $[75 - 100]$ , which we denote as *low*, *medium*, *high*, and *very high*, respectively. This makes the individual scores more intuitive than actual percentages. The higher the values are, typically the more negative is the visual impact. However, in a use case where urban planners want to create a new landmark building they might want to choose a candidate with a high or very high candidate visibility. We additionally use the ‘all candidates’ image to compute a ratio describing the relevance of a viewpoint. For this we divide the number of pixels covered by any of the candidates by the total number of image pixels. This value indicates how much of the candidate geometry is actually visible from this viewpoint and implies how important visual-impact values in this viewpoint are. Since these percentages are even more difficult to interpret, we bin them into three classes as depicted in Figure 3.2e:  $[0 - 5[$ ,  $[5 - 10[$ , and  $[10 - 100]$ , corresponding to *small*, *medium*, and *large* coverage. We designed this classification empirically, based on feedback from our collaborators.

### 3.7 Use Case Scenario

At this point, we introduce a use case scenario, to make the discussion of design decisions in Section 3.8 and the subsequent demonstration of the applicability of Vis-A-Ware in Section 3.9 more tangible. In cooperation with the Vienna department of urban surveying, we obtained a detailed 3D city model of Vienna and of the underlying terrain. Since the rights for each submitted

### 3. VIS-A-WARE: INTEGRATING GEOMETRIC AND ATTRIBUTE VIEWS FOR VISIBILITY-AWARE URBAN PLANNING




Figure 3.3: 20 exemplary candidate buildings of different sizes and shapes identified by letters from A to T.

candidate building lie with the respective architecture firm, it was not possible to use data from a real call for proposals, therefore this is an artificial use case. As illustrated in Figure 3.3, we created 20 candidate buildings of various shapes and sizes in accordance with our experts and placed them in a project area near Vienna's north railway station (Figure 3.4). Each candidate is identified by a letter from *A* to *T*. In the traditional workflow, expensive computations limit the analysis to a handful of viewpoints. However, domain experts stated that using more viewpoints would be desirable, for instance, along certain streets or railroad tracks. Since this restriction does not apply to *Vis-A-Ware*, we placed a total of 100 viewpoints in relevant areas of the city, identified by an id from *vp1* to *vp100*, of which 53 are in the vicinity of the planning area. We divided them into three categories: 'LoS' (7) being lines of sight, 'Street' (32) being viewpoints on street level, and 'Train' (14) being viewpoints along the elevated train tracks going past the project area. The scene further includes one landmark, the St. Stephen's Cathedral, colored in orange, as shown in Figure 3.4.



Figure 3.4: (a) Bird's-eye and (b) close up view, both showing the project area near the railway station with all building candidates and a landmark (orange). Further, we show exemplary viewpoints at street level and on the train tracks, as well as a line of sight.




Figure 3.5: An overview of Vis-A-Ware, a tool supporting 3D visibility analysis in urban planning. It features the following coordinated views: (a) a 3D spatial view to investigate spatial relations between planned structures and viewpoints, (b) the transposable ranking view to analyze derived visibility metrics, and the filmstrip (c) for a qualitative comparison of single-building renderings. (d) 3D representations of the candidate buildings are blended together, hovering in linked views renders the respective candidate as solid.



## 3.8 Visualization and Interaction Design

Vis-A-Ware consists of three coordinated views: The *transposable ranking view (TRV)* (Figure 3.5b) is the center piece of Vis-A-Ware and allows users to interactively explore all quantitative evaluation data. The *3D spatial view* (Figure 3.5a) offers an interactive exploration of the cityscape, candidates, and viewpoints. Finally, the *filmstrip* (Figure 3.5c) provides a side-by-side comparison of candidate renderings from specific viewpoints for a qualitative evaluation.

### 3.8.1 Transposable Ranking View (TRV)

The TRV was designed with the following goals in mind:

- **Overview** of visual impact results concerning candidates and viewpoints (**G2.4**).
- Intuitive **filtering** of candidates and viewpoints (**G2.5**).
- **Detail-on-demand** showing individual visual impact values (**G2.1**).
- **Ranking** of candidates or viewpoints by individual VIMs (**G2.3, G2.5**).
- **Comparison** of candidates by multiple VIMs (**G2.2, G2.5**).

The TRV is a tabular view, where the rows either represent candidates or viewpoints (confer with lookup table in Figure 3.1), while the individual VIMs determine the columns of this table. Each cell contains the distribution of the minor items regarding a specific VIM. For instance, Figure 3.6 shows VIM values of candidates for each viewpoint, since the TRV is in viewpoint-major mode. Indicated by the bar charts, the three cells shown in Figure 3.6b contain candidate visibility, openness occlusion, and sky occlusion values for each candidate.

#### Visual Encoding

The TRV in Figure 3.6 depicts the evaluation data in viewpoint-major mode showing one row for each viewpoint. The *detail bar charts* in Figure 3.6a show the VIM scores of each candidate in a single viewpoint. The detail bars are sorted by candidate score (i.e., bar height) and encode the respective class into the color of a bar. Through the color coding and the sorting we can condense the bar charts to *stacked bar charts* as depicted in Figure 3.6b, which allows us to efficiently display a large number of rows. If hovering a row, it expands and reveals the detail bar charts providing detail-on-demand. Hovering individual detail bars opens a tool tip, which shows the exact VIM value of a candidate (Figure 3.6c). This detail hovering further triggers a peek-brush [BP10b] interaction, which shows the occurrence of the respective candidate in all other VIM distributions by showing the candidate's letter in the stacked bar chart. As illustrated in Figure 3.6b, the rank of a candidate within a distribution is preserved which effectively conveys the ranks of a candidate for each VIM and each viewpoint.

We used colorbrewer [HB03] to pick color schemes for visually encoding the classes of the different VIMs. For candidate visibility, openness occlusion, and sky occlusion, we chose




Figure 3.6: Transposable Ranking View (TRV). (a) Detail bar charts showing individual VIM values of each candidate. (b) Stacked bar charts as a compact version of VIM distributions. (c) Linked peek-brushing shows detail-on-demand tooltip and a candidate's position within all other VIM distributions. (d) Rows are ranked by distribution scores of the selected VIM - in this case candidate visibility. (e) The arrow-icon loads the perspective from a viewpoint into the 3D view, while the film icon loads respective screenshot data into the filmstrip. (f) Each category of viewpoints is collapsible and features a heatmap-like aggregation of VIM scores. (g) Viewpoints are grouped by their coverage value, which is indicated by the size of the black circles.

sequential 4-class single hue color schemes, with the respective hues being red, green, and blue. Since the landmark evaluation is a more decisive VIM and slightly differs in computation from the others, we chose a multi-hue color scheme from yellow to red. The color coding can be adapted and easily replaced with other colorbrewer schemes.

### Focus, Filter, and Transpose

The workflow supported by the TRV is a sequence of *focus*, *filter*, and *transpose* interactions, which allow users to investigate evaluations with respect to candidates and viewpoints. When clicking on individual stacked bars (Figure 3.7a) the respective minor items are added to the focus set while the others are considered as context. To visually distinguish focus and context we create *split stacked bars*. As illustrated in Figure 3.7b, the focus on the left is rendered at full height, while the stacked bars of the context on the right are rendered at reduced height. This partition of minor items into focus and context is reflected in all stacked bar charts. An arbitrary concatenation of focus interactions allows users to create a focus set by selecting any number of stacked bars (Figure 3.7c). Through keyboard interaction, the minor items of the context set are filtered from the TRV (Figure 3.7d). The transpose interaction allows users to shift the focus of the analysis from viewpoints to candidates or vice versa by switching between candidate-major and viewpoint-major mode. For instance, Figure 3.7d shows viewpoints as rows (viewpoint-major), while after transposing Figure 3.7e shows candidates as rows (candidate-major).

For demonstration purposes Figure 3.7 illustrates the identification of candidates that are most prominent as seen from the railway (compare Section 3.9.2) by depicting relevant cutouts of the TRV for each interaction (Figure 3.7a-i). The displayed data set comprises twenty candidates and six relevant viewpoints from the ‘Train’ category. To focus on candidates, users are able to select individual stacked bars. For example, selecting the high visibility scores of viewpoint *vp50* in Figure 3.7a, splits the stacked bar charts into two charts as shown in Figure 3.7b. The split stacked bars allow users to judge how the candidates in focus impact other viewpoints. Users may select additional high class bars among the context candidates (Figure 3.7b), which are then added to the focus set (Figure 3.7c). All high impact classes are in focus, so users filter the context candidates (Figure 3.7c). This expands the focus set to cover the whole width of the chart (Figure 3.7d). For reducing the set of viewpoints, transposing the TRV makes the candidates the major items, while the viewpoints are encoded into the stacked bars (Figure 3.7e). This enables the selection of viewpoints containing a high visibility value for candidate *T* (Figure 3.7e). Figure 3.7f indicates that other viewpoints exhibit high impact values from other candidates as well, which can be included into the selection. All viewpoints with high classes are in focus (Figure 3.7g). Filtering of the context leaves six candidates and four viewpoints (Figure 3.7h and Figure 3.7i).

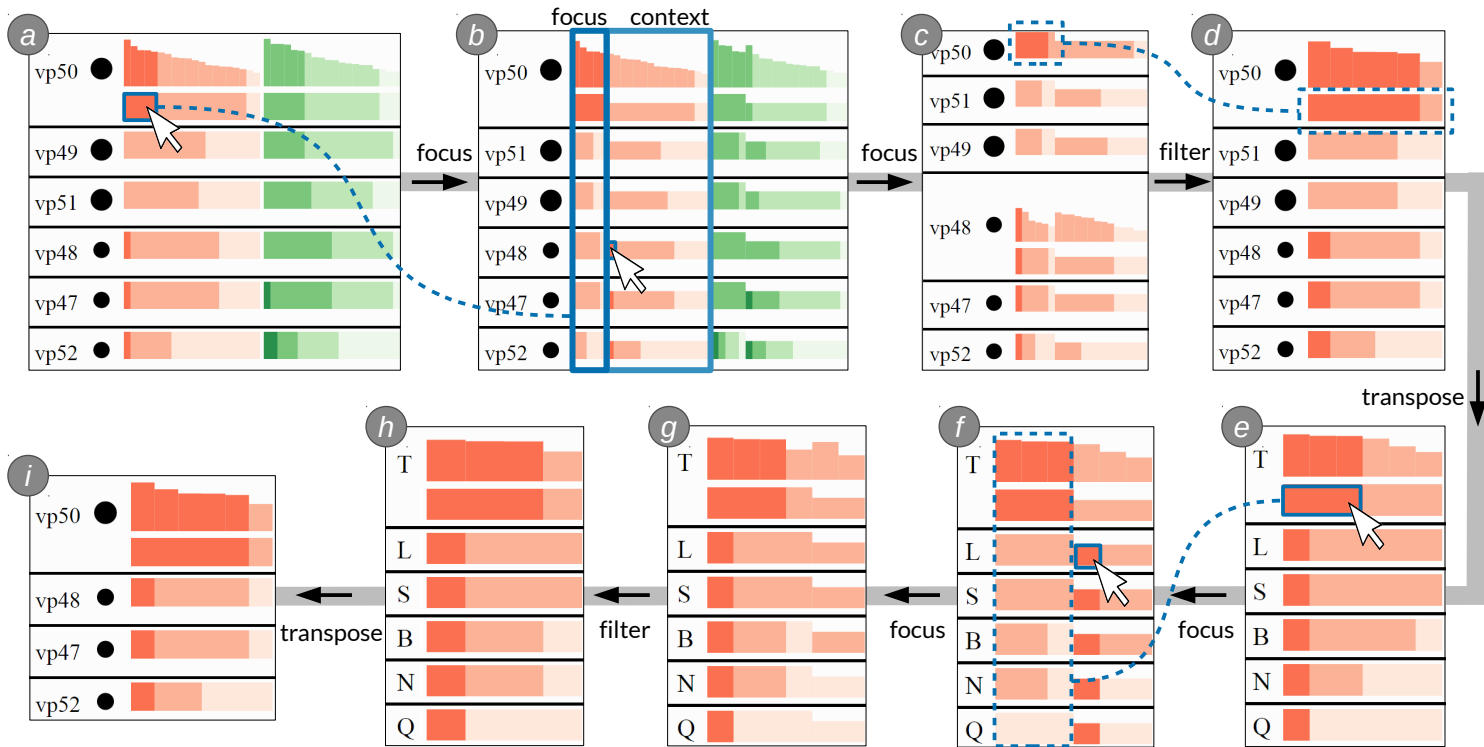


Figure 3.7: Focus, Filter, and Transpose workflow for reducing candidate and viewpoint sets. (a) Twenty candidates impact six viewpoints. Selection of high class visibility candidates in viewpoint *vp50*. (b) Higher bars encode selected candidates (in focus). Selection of additional high class visibility candidates in viewpoint *vp48*. (c) The focus set covers all candidates exhibiting high visibility values. (d) Candidates are filtered to the focus set, dropping all candidates, which do not exhibit high visibility values in any viewpoint. (e) Transposed view: candidates are now rows, the selection now affects viewpoints highly impacted by candidate *T*. (f) Selection is extended by other highly impacted viewpoints. (g) All highly impacted viewpoints are now part of the focus set. (h) Set is filtered to the focus set. (i) Transpose to see the set reduced to six candidates in four viewpoints.

### Ranking

To quickly identify highest scoring candidates or highly impacted viewpoints, the rows of the TRV must be sortable by any VIM. Therefore, we need to assign an unambiguous score to any distribution of visual impact. This score has to ensure that, for instance, a viewpoint exhibiting a single candidate in the very high class is always ranked higher than any viewpoint not containing a very high VIM score. Thus, we use the mathematical concept of a numeral system to assign a score to each stacked bar chart:

$$\text{score}(d) = \sum_{i=1}^p x_i \cdot b^{i-1} \quad (3.3)$$

where  $d$  is the distribution of VIM values,  $p$  is the number of classes used in  $d$ ,  $b$  is the maximum number of items in one class (total number of viewpoints or candidates) and  $x_i$  is the actual number of items in class  $i$ . As it can be seen in Figure 3.6 and Figure 3.7, viewpoints are ordered within their category depending on their candidate visibility values. In terms of a numeral system,  $b$  refers to the base while  $p$  represents the number of digits.

### Categories, Aggregation, and Grouping

Not all viewpoints in a city have the same semantic meaning, for instance, lines of sight, street-level viewpoints, and viewpoints along a railway section as in the presented use case scenario. Vis-A-Ware thus supports to divide the set of viewpoints into user-defined categories. Each category is represented by a collapsible table while each category header features an aggregation bar chart per VIM (Figure 3.6f). The aggregation bar charts show a vertical stacked bar for each candidate, while each bar represents the number of viewpoints per VIM class.

This transposed aggregation allows users to quickly identify candidates, which exhibit high values in a category. Furthermore, it presents a heat-map-like overview of VIM distributions and indicates if a certain category is worth exploring. For the target use case of deciding between individual candidates, only viewpoints are reasonably divided into semantic groups, and therefore the collapsible categories are only present in viewpoint-major mode. Within each category, viewpoints are further grouped by their coverage value (Figure 3.2e) encoded by a black circle next to the viewpoint identifier as illustrated in Figure 3.6g. Three different circle sizes represent the three different coverage classes large, medium, and small. If a viewpoint category is collapsed its viewpoints are excluded from the TRV. Consequently, if transposing to candidate-major mode, these viewpoints do not appear in the stacked bar charts.

### Integrated Interaction

The peek-selection interaction through hovering, as discussed previously in this section, allows users to quickly highlight candidates or viewpoints in all charts (Figure 3.6c). The TRV in viewpoint-major mode features an arrow-icon (Figure 3.6e), which if clicked changes the perspective of the 3D geometric view to match the perspective of the selected viewpoint. Clicking the film icon loads the associated data row into the filmstrip for a qualitative evaluation (Figure 3.5c). The ranking, the filtering, and the peek-selection are coordinated with the 3D spatial view and

the filmstrip. How these views react to these interactions, will be explained in Section 3.8.2 and Section 3.8.3, respectively.

### 3.8.2 3D Geometric View

The 3D geometric view depicts a 3D geospatial environment, consisting of a city model, landmarks, candidate buildings, and circular glyphs representing the positions of viewpoints. Candidates are rendered as semi-transparent colored 3D models, which enables the users to inspect multiple buildings at once. Since it is not possible to pick distinguishable colors for about 30 candidate buildings, we chose a set of 10 repeating qualitative colors [HB03]. The set of displayed candidates can be reduced through filtering in the TRV. Peek-brushing of candidate items in the TRV or the filmstrip causes 3D representations of candidates in the geometric view to be rendered completely opaque, as illustrated in Figure 3.5d. This enables users to compare the spatial properties of one particular candidate with respect to the other candidates of the currently displayed set.

As it is depicted in Figure 3.5a, we display an aggregated view of the candidate distribution in each viewpoint as circular glyphs. A glyph comprises a colored circle containing a number and is displayed at the spatial position of its respective viewpoint. This enables users to assess the spatial distribution of VIM values across the city scene. If major items in the TRV are ranked by a specific VIM, this VIM is encoded by the color of the circles. The specific color of the circle corresponds to the highest class in its distribution, while the number in it is the count of candidates in this respective class. The coverage value of a viewpoint is mapped to the area of the circle. This enables users to judge the extent of the visual impact, with respect to the selected VIM, across the city. Viewpoints with a small coverage are de-emphasized. Only the viewpoints of expanded categories in the TRV are shown in the described encoding. All other viewpoints are considered as context and are rendered as gray circles. Peek-brushing of viewpoints in other views causes the highlighting of the respective circles in the geometric view by emphasizing their border line.

### 3.8.3 Filmstrip

The domain experts are accustomed to side-by-side comparisons for the qualitative analysis of candidates. We chose a filmstrip metaphor, similar to Bruckner et al. [BM10], to display the images which are associated with each evaluation in a side-by-side layout, as depicted in Figure 3.5c. Clicking the film icon in the TRV loads the content of the respective row into the filmstrip. Depending on the current mode of the TRV, the filmstrip either shows the images of all candidates from one viewpoint (viewpoint-major) or one candidate from all viewpoints (candidate-major). A header line shows the identifier of the major item, i.e., the row, while a rectangle in the upper left corner of each image displays the name and VIM value of each minor item. An image's border color encodes the impact class of the corresponding candidate-viewpoint pair. The displayed VIM corresponds to the VIM currently selected for ranking in the TRV. The images displayed in the filmstrip are filtered depending on the filter criteria specified in the TRV. Hovering images causes highlighting of the respective items in the other views.




Figure 3.8: Q1. Which candidates do cover a landmark and how strong is the occlusion? (a) Users rank by landmark occlusion and (b) focus on very high and high impact values. (c) All very high and high classes are part of the focus set and (d) filtering reduces the set of candidates to five. (e) To investigate top scoring candidates *F* and *T* in detail, the TRV is transposed to candidate-major mode. (f) Users focus on viewpoints exhibiting very high impact values, which are then assessed through (g) a qualitative comparison in the filmstrip.

### 3.8.4 Implementation

For the interactive rendering of the 3D urban scene we use the Aardvark real-time rendering engine [VRVa] written in F# [MR]. The TRV and the filmstrip are HTML5 pages displayed in a chrome browser, utilizing d3.js [BOH11]. The viewpoint circles are also rendered to an HTML5 page, which is integrated as an overlay into the 3D geometric view with the help of CEF (Chrome Embedded Framework) [Gre]. Both parts, the F# application and the HTML5 views, are communicating by exchanging JSON messages over a WebSocket communication layer.

## 3.9 Visibility Analysis

In collaboration with domain experts from urban planning we formulated three exemplary questions based on the traditional qualitative visibility-analysis workflow. We demonstrate how the developed views and metrics support the users to navigate through evaluation results and inspect selected subsets of candidates and viewpoints qualitatively.

### 3.9.1 Which candidates cover a landmark and how strong is the occlusion? (Q1)

Severely occluding a landmark along a line of sight is an exclusion criterion for candidates. To focus on lines of sight users collapse all categories except the ‘LoS’ one. The remaining viewpoints are ranked by the landmark occlusion metric as shown in Figure 3.8a. Hovering the aggregation view reveals that *F* and *T* cause very high landmark occlusions (Figure 3.8a). Apparently only the lines of sight *vp5* and *vp4* are affected, which is indicated by the absence of orange bars in *vp6* and *vp3*. To focus on very high and high classes of landmark occlusion, users are able to select the respectively colored bins in *vp4*. The split stacked bars in Figure 3.8b show that other candidates than the ones previously selected cause very high and high occlusions in *vp5*. Users select the respective candidates in *vp5* and filter the TRV to the selected set of candidates (Figure 3.8c). Inspection of the detail bar charts uncovers, for instance, that *F* has a very high score in *vp4*, but only medium impact in *vp5* (Figure 3.8d). To analyze the individual impact values of the selected candidates with respect to *vp4* and *vp5* analysts transpose the TRV (candidate-major) (Figure 3.8e) and focus and filter so that only viewpoints *vp4* and *vp5* remain (Figure 3.8f). The data set is now reduced to five candidates in two viewpoints. Analysts can now evaluate each candidate qualitatively by opening the filmstrip, depicting the selected candidate as seen from two viewpoints. As Figure 3.8g shows, the landmark occlusions of *F* in *vp4* and *T* in *vp5* are indeed very high.



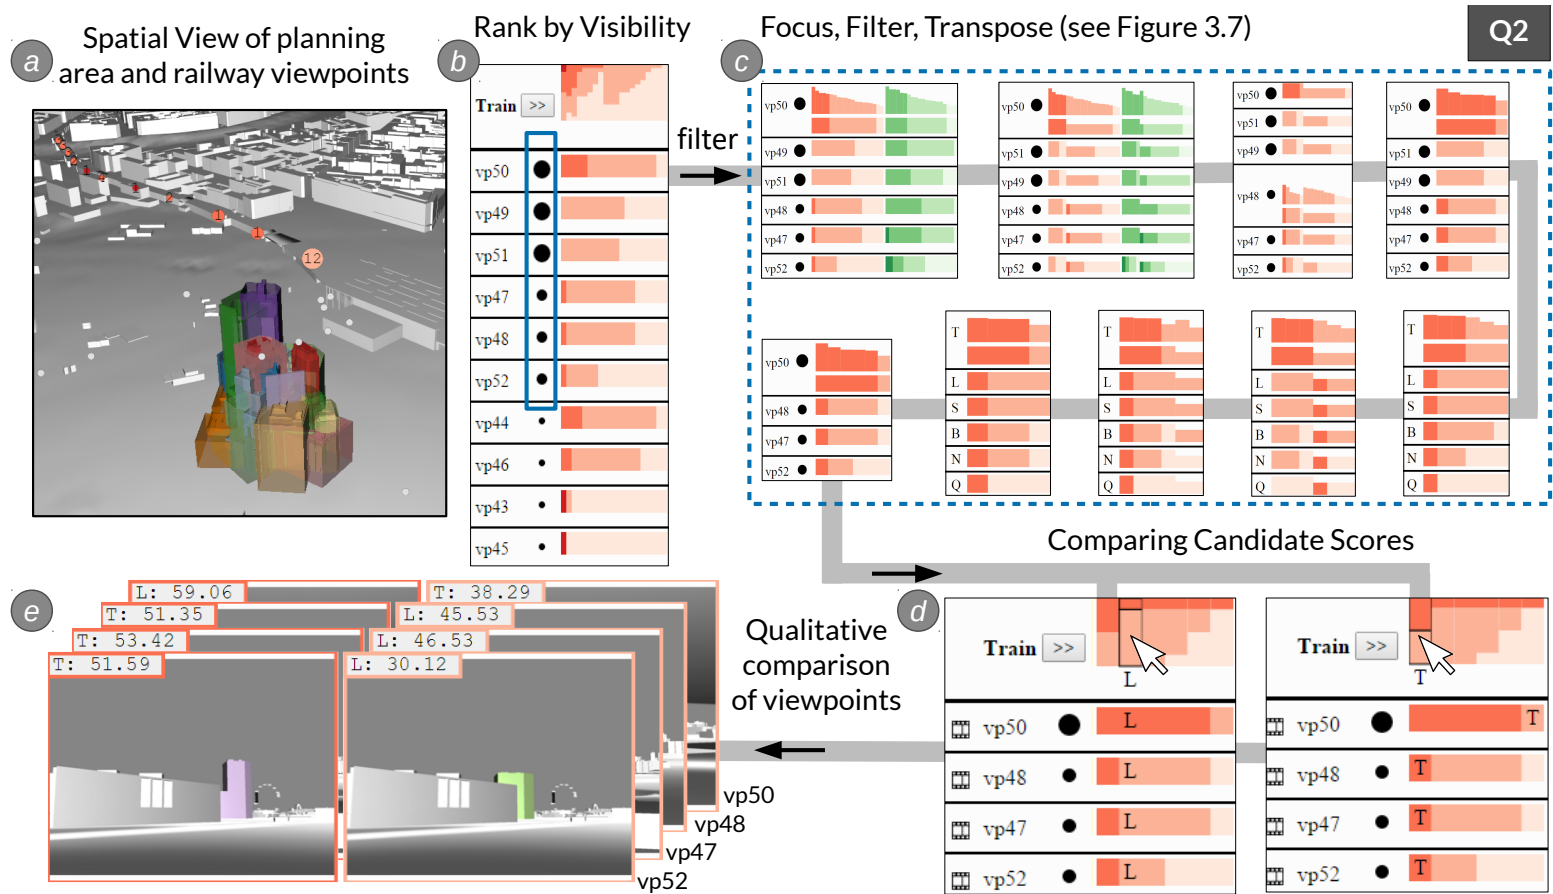


Figure 3.9: Q2. Which candidates are most prominent as seen from the railway? (a) A spatial overview depicts the distribution of the visibility impact as seen from the railway. (b) Viewpoints are sorted by candidate visibility. (c) Candidates and viewpoints are reduced through multiple focus, filter, and transpose interactions as illustrated in Figure 3.7. (d) Peek-brushing of the highest scoring candidates reveals their scores within the distributions. (e) The remaining candidate-viewpoint combinations can be compared qualitatively in the filmstrip.

### 3.9.2 Which candidates are most prominent as seen from the railway? (Q2)

High visibility is not necessarily a negative aspect of a candidate. Often city administrations want to construct a building with a certain prominence and recognition value. To find the most prominent candidates along the railway users expand the ‘Train’ category in the TRV and rank the viewpoints by visibility. The 3D spatial view shows the distribution of affected viewpoints on the railroad (Figure 3.9a). The sizes of the circles indicate that many viewpoints exhibit only a small coverage ratio, which deems these viewpoints as rather irrelevant for the question at hand. This is also reflected by the black circle in each viewpoint row as depicted in Figure 3.9b. The aggregation shows multiple very high occurrences. The stacked bars, however, reveal that these only occur in viewpoints with small coverage (*vp43*, *vp45*). Users consider viewpoints of large and medium coverage only and filter the TRV accordingly. Then they apply a series of focus, filter, and transpose interactions to reduce the set of viewpoints and candidates as shown in Figure 3.9c. This is discussed in detail in Section 3.8.1 and Figure 3.7. The interaction sequence begins with the reduced viewpoint set depicted in Figure 3.7a. The aggregation view shown in Figure 3.9d indicates that candidate *T* is most prominent as seen from the train tracks. However, peek-brushing the aggregation bars reveals vastly different distributions of the candidates *T* and *L*. Candidate *T* only has a medium score in the large coverage viewpoint *vp50*, while candidate *L* scores much higher in *vp50* and relatively high in the medium coverage viewpoints. Which candidate actually is considered to be more prominent, is subject to discussion supported by a qualitative comparison in the filmstrip.

### 3.9.3 Which candidates have the highest impact on sky and openness in the vicinity of a project area? (Q3)

In this use case, as in the previous section, only viewpoints with a large or medium coverage value are considered. First, users sort the set of the 24 remaining viewpoints by openness occlusion, which reveals that for *vp63* all candidates have a very high openness score. Since this viewpoint is not very decisive, users select very high classes in the other viewpoints (Figure 3.10a and Figure 3.10b). Then they sort the viewpoints by sky occlusion, and select candidates with high sky occlusion values (Figure 3.10c and Figure 3.10d), and filter to the selected candidate subset (Figure 3.10e). Transposing the TRV depicts the remaining six candidates and allows users to select viewpoints, which exhibit very high and high occurrences of openness occlusion and sky occlusion, respectively (Figure 3.10f and Figure 3.10g). After filtering to the selected set (Figure 3.10h) users explore the highest scoring candidates through sorting by the respective VIM. As shown in Figure 3.10i candidate *O* has the highest openness occlusion while candidate *T* causes the highest sky occlusion. If hovering the candidate row the candidate building is highlighted in the spatial view. While further hovering individual bars, the respective viewpoint glyphs are highlighted, which allows users to assess the spatial distribution of impacted viewpoints.




Figure 3.10: Q3: Which candidates have the highest impact on sky and openness in the vicinity of the project area? (a) Users rank viewpoints by the openness metric and focus on (i.e., select) very high values of openness occlusion. (b) The focus set now contains all very high openness values ignoring *vp63*. (c) Then users rank by sky occlusion and select high values. (d) The focus set contains all high sky occlusion values. (e) Now users filter candidates, which are not in focus. (f) Transposing the view lets the users investigate the remaining six candidates. (g) They again focus on viewpoints exhibiting very high openness occlusion and high sky occlusion and (h) filter to the focus set. (i) Sorting by either openness occlusion or sky occlusion yields the top ranked candidates *O* and *T* respectively. Hovering candidate rows and viewpoints shows detail bar charts and results in the highlighting of buildings and viewpoint circles in the spatial view.

### 3.9.4 User Feedback

During the design phase we collected feedback regarding prototype iterations of Vis-A-Ware with a group of ten participants from the departments of 'Urban Surveying' and 'District Planning and Area Usage' of Vienna. Additionally, we discussed design decisions and VIMs with an urban planning expert. For the final evaluation we interviewed six domain experts from three different fields concerned with 3D visibility analysis: urban surveying (A), urban planning (B), and architecture and urban design (C). Each field was represented by two experts. Experts from A and B typically collaborate to create 3D visibility analysis products, while group C participates in the actual decision-making. In general, the feedback from these sessions was very positive. Following a pre-defined protocol, each evaluation started with an overview of the system to reintroduce the individual components and their interplay, followed by a recapitulation of the VIM computations. After an average period of twenty minutes for familiarizing themselves by actively using ranking, focus, filter, and transpose, the experts found the interactions comprehensible and intuitive. Given the possibility for questions and support they could work through the three use cases presented in Section 3.9.1, Section 3.9.2, and Section 3.9.3.

Users from A and B could immediately see that an interactive tool like Vis-A-Ware would improve their workflow by facilitating the evaluation of a large number of viewpoints along streets. Further, they stated that our approach for line-of-sight evaluation and landmark occlusion is a very fast and intuitive alternative to time-consuming 2D visibility maps computed by GIS applications. They envisioned to use Vis-A-Ware as their primary tool for collecting and quantifying visibility data. They welcomed the integration of an objective measure of visibility (i.e., VIMs) with the qualitative analysis provided by the filmstrip and the 3D spatial view.

The direct visual linking between spatial distribution and VIM values was appreciated as well, because it allows users to judge which parts of the city are impacted on a granular level. Users from A and B pointed out, that Vis-A-Ware could efficiently evaluate whole streets, which has been a tedious task in the past.

When trying to comprehend the meaning of individual metrics, users of group A and B sometimes had difficulties to interpret the rather abstract VIM values for candidate visibility, openness occlusion, and sky occlusion. They stated that they would benefit from an attribute indicating the shape of a candidate. Two candidates, one tall, one wide, could essentially have the same sky occlusion value, but both candidates would have a different perceived impact on the cityscape. Furthermore, they questioned the relevance of the openness occlusion metric, since investigations about openness occlusion are typically not part of their analysis. They were very fond of the implementation of the landmark occlusion metric, because it allowed them to quickly identify viewpoints from where the cathedral is occluded. They were able to judge the degree of occlusion quantitatively through the bar charts and verify it qualitatively through the images in the filmstrip.

Decision-makers (experts from group C) are not the primary user group of the presented tool, since only a selected subset of viewpoints is presented to them for qualitative comparison. However, they could see the benefits of using the spatial view of Vis-A-Ware in jury discussions, where also project neighbors and political decision-makers are involved. Discussing a miniature model of the cityscape and candidates, the decision-makers can only assess the visual impact from a bird's eye

viewpoint. This issue is currently alleviated by placing a miniature camera into the haptic model to emulate the perspective of a pedestrian. A 3D visualization as provided through Vis-A-Ware is more flexible and viewpoints can be changed interactively. Furthermore, the experts from group C were very interested in exploiting the interactivity of the tool. They suggested to perform and visualize evaluations along paths, such as signature streets or tourist walks, with different movement profiles (e.g., pedestrian, bicycle, or car movements) and playing back animations along these paths. They would also like to employ Vis-A-Ware to evaluate an upcoming high profile project regarding a new museum.

The user feedback tells us that Vis-A-Ware makes visibility analysis for experts from groups A and B more holistic, more objective, and less time-consuming. For experts from group C, Vis-A-Ware presents opportunities to evaluate, analyze, and visualize dynamic movement paths.

### 3.9.5 Performance

We use a system comprising an Intel i7-3770 quad-core CPU running at 3.40 GHz, 16 GB of RAM, and a GeForce GTX 680 graphics card. The 3D rendering runs smoothly at 50-70 frames per second, while rendering 500k of triangles for buildings and 6000k faces for the terrain. We processed 20 candidate buildings against 106 viewpoints resulting in 2332 ( $= (20 + 2) * 106$ ) evaluations including the existing situation and the ‘all candidates’ images. For the VIM computations, we rendered 2332 false-color images of reduced size (640x480 pixels), which took approximately 50 seconds. Taking 2120 screenshots (1280x960 pixels), one for each viewpoint-candidate pair, for a qualitative analysis took approximately 277 seconds. The smaller false-color images are processed in-memory, while the larger screenshots are saved to the hard disk, which is the most expensive step of the evaluation. With the given system and quality parameters we can process 7.1 ( $= 2332 / (277 + 50)$ ) candidate-viewpoint pairs per second.

## 3.10 Summary

At this point we summarize the key aspects of Vis-A-Ware and relate them to the design goals formulated in Section 3.2. We defined four metrics to quantify the visual impact of candidate buildings - landmark occlusion, candidate visibility, openness occlusion, and sky occlusion (**G2.1**). The transposable ranking view (TRV) visualizes distributions of these metrics with respect to candidates and viewpoints. The distribution visualization provides a detail view on individual evaluation values (**G2.3**). The TRV provides interactions for ranking, focusing, filtering, and transposing of the data set (**G2.3**, **G2.5**, **G2.6**). Contents of a TRV row can at all times be loaded into the filmstrip, which enables users to access image data for a qualitative analysis (**G2.2**, **G2.6**). Linked peek-selection connects items in all views by hovering and highlighting (**G2.2**, **G2.4**). The geometric view allows users to interactively explore the scene, which provides spatial context for candidates and viewpoints including their evaluation values.

### 3.10.1 Visual Scalability

The TRV in viewpoint-major mode provides a clear overview of viewpoints and scores of dozens of candidates through aggregation. The stacked bar chart representation is designed to efficiently use vertical space. Through categorization of viewpoints and filtering of viewpoints with low coverage, the number of viewpoints considered at once can be reduced significantly. Alternating filter and transpose interactions enable users to break down the problem space into a small set of items to be inspected in detail. The current implementation of the TRV supports four different metrics. However, it could be extended to an arbitrary number of metrics, by aggregating, resizing, or collapsing of columns as presented in Gratzl et al. [GLG<sup>+</sup>13]. The underlying data model already supports an arbitrary number of metrics per evaluation.

### 3.10.2 Applicability and Generalization

Vis-A-Ware features several linked views for a quantitative and a qualitative analysis of multiple visual impact metrics of candidates with respect to viewpoints. This approach directly translates to many visibility-related decision problems commonly solved with GIS applications. For instance, in military planning different path options can be evaluated for their exposure with respect to enemy positions and vantage points.

We can abstract our approach from candidates, viewpoints, and visibility analysis to the more general terms of discrete variations, probes, and simulation. Applied to a disaster management use case, we could consider a flooding simulation, where different barricading plans exist as discrete variations and important buildings act as probes for damage values. Users could easily identify bad barricading plans and buildings prone to flooding.

The TRV, as the center piece of this design study, can be used for ranking alternatives in any multi-criteria decision problem, such as, building design parameters [DFL<sup>+</sup>15] or sites for urban development [FLD<sup>+</sup>15]. It could also help to explore a large number of public utility issues in the context of districts and responsible departments [ZYM<sup>+</sup>14]. The TRV is designed to explore a two-dimensional matrix, where each cell contains a vector of arbitrary length. It could support biologists in studying gene expressions, where they typically have to explore multiple potentially large heatmaps.

### 3.10.3 Limitations and Future work

Expert feedback showed that Vis-A-Ware has clear benefits over traditional, purely qualitative workflows. However, based on feedback sessions and design discussions we could identify points of improvement and opportunities for new features as future work.

- Whereas experts could easily relate to landmark occlusion, candidate visibility, and sky occlusion, the meaning of the openness occlusion was deemed as rather unintuitive. We believe this lies in the volumetric nature of openness perception as described by Yang et al. [YPL07]. Extending our evaluation approach by additionally comparing depth images could yield more meaningful results and could actually quantify the volume of the open space occluded by a candidate.

- The 3D visibility analysis is a crucial, but actually small part in the decision-making process. There are many non-visual parameters, such as expected revenue or number of apartments and offices, which factor into a final decision. Extending the columns of the TRV by non-visual parameters or scores from traditional automatic MCDA algorithms could make Vis-A-Ware a holistic tool for decision-making in urban planning.
- Candidate buildings potentially cast unwelcome shadows onto public squares or neighboring houses. While the TRV could be extended by a derived shadow VIM, the computation of that value and the qualitative comparison of cast shadows in the geometric view over time are potentially challenging.
- Feedback from experts (in particular group C - architecture and urban design) indicated a need for evaluating visibility along paths with respect to different movement profiles. The current evaluation approach of Vis-A-Ware is suitable for computing VIMs along pre-defined paths. Designing a visualization for analyzing multiple visual impact metrics of 30 candidates over a spatio-temporal domain (movement profiles) poses an interesting challenge.





# InCorr: Interactive Data-Driven Correlation Panels for Digital Outcrop Analysis

**This chapter is based on the following publication:**

*Thomas Ortner, Andreas Walch, Rebecca Nowak, Robert Barnes, Thomas Höllt, and Eduard Gröller. InCorr: Interactive Data-Driven Correlation Panels for Digital Outcrop Analysis. In Proceedings of IEEE Transactions on Visualization and Computer Graphics, 27(2):755–764, February 2021*

Geological analysis of image data collected by stereo camera systems on Mars rovers has proven to be a valuable tool in reconstructing ancient environments on Mars. This is a key aspect of the upcoming ESA ExoMars 2022 *Rosalind Franklin* Rover and the NASA 2022 Rover *Perseverance* missions [ESA, NAS]. A shared aim of both of these missions is to drill and sample rocks, which were deposited in ancient environments that scientists deem may have been habitable. This is determined largely by the geological characteristics. Within geology, the field of sedimentology is concerned with the analysis of textures and internal fabrics of rocks formed by deposition and movement of sand, silt, and clay sediment by wind, water, or ice. Burial and exposure to high heat, pressure, and fluid circulation leads to the formation of sedimentary rocks. Stratigraphy is a field concerned with documenting and interpreting the vertical distribution of different textures and sedimentary structures in order to reconstruct the evolution of the environments, which formed the sedimentary layers, or *strata*. There is strong scientific evidence that there were rivers and lakes active at the surface in a distant past [GGM<sup>+</sup>15]. Assuming analogous processes on Mars as on Earth, lake and river deposits are most promising for scientists to discover biosignatures. Therefore, stratigraphic analysis and correlation of observations between distant locations is essential.

## 4.1 Remote Stratigraphic Analysis

The main goal of geologists is to combine their observations to build a geological model, which encompasses the temporal evolution of environmental processes of a region. They build such a model from meticulously annotating and measuring textural features in a number of *outcrops*, i.e., rock faces exposing strata as illustrated in Figure 4.1, and determining their relative age-relationships. It is generally accepted that younger sediments are deposited on top of older sediments, therefore changes in rock characteristics are a record of its past. Correlation of these observations across large areas allows for regional evolutionary models to be built. Traditionally, geologists make measurements ‘in the field’ using hands-on tools, such as a compass clinometer to measure surface orientation. The development of affordable remote sensing solutions has prompted a sharp rise in geological analyses of *digital outcrop models (DOM)*. DOMs present the geologists with 3D triangulated, and often photographically textured, surfaces of rock outcrops. Application of these techniques to processed stereo-camera image data collected by rovers on Mars can be used to greatly enhance our understanding of the evolution of the planet and drive future robotic exploration missions.

Robust and efficient interpretations and measurements can be collected from DOMs in PRo3D [BGT<sup>+</sup>18]. These data are the foundation for creating a regional geological model. Outcrops can be understood as sampling locations, partially exposing a succession of strata, which in fact may extend over hundred thousands of square kilometers underneath a planet’s surface. To build a regional geological model, geologists look for occurrences of the same stratum in multiple outcrops. First, they characterize the succession of exposed strata from an outcrop by creating a *geological log*, as shown in Figure 4.6a. After repeating this for each outcrop, they combine all logs into a *correlation panel*, and connect matching strata to form *correlations*, as indicated by the red lines in Figure 4.6b. Correlating strata across multiple outcrops that cover a large area leads to a reliable characterization of the distribution of important geological units.

In a long lasting cooperation with planetary geologists, who are engaged in the geological survey of Mars as part of the scientific working groups of ESA and NASA missions, we identified the following challenges that come with a remote geological analysis. At the moment, DOM visualization software is tailored towards the *interpretation* of outcrops, i.e., annotating and measuring strata. For further analysis, geologists typically rely on the export of measured values to tools for statistical analysis and graph plotting, while they create correlation panels manually based on the obtained values. Geologists often use vector drawing software or they might draw the panels completely by hand. In any case, the creation of correlation panels is very time-consuming and therefore it is typically left to the end of the workflow. However, only then potential weaknesses in the interpretation data become apparent, such as, insufficient detail in the interpretation of a particular outcrop. This requires a tedious revisit of the interpretation stage and forces geologists to reorganize and largely redraw their panels. Further, the separation between digital tools enforces a separation of interpreting and correlating, which would not be the case in the workflow of the traditional field work. Ultimately, this separation disconnects correlation panel illustrations from the underlying annotations and measurements, which inhibits traceability, reproducibility, and reusability of analysis results throughout the geological science community.




Figure 4.1: (a) Map of the Hanksville-Burpee Dinosaur Quarry campaign with its four outcrops along the canyon. (b) 3D triangulated mesh as digital outcrop model (DOM) of Sol2 with interpretation showing bedsets and units and their respective contacts. (c) Cross bed measurements, which determine a layer's deposition direction indicated by dip-and-strike disks.

In this chapter, we present the visualization solution *InCorr*, short for *Interactive data-driven Correlations*. Based on the long lasting cooperation with leading scientists in the field of planetary geology, we conducted a design study to create a 3D geological logging tool and the *InCorrPanel*, an interactive data-driven correlation panel. Both components are integrated into PRo3D [BGT<sup>+</sup>18], a tool for the geological interpretation of DOMs. We verify the applicability of InCorr through a use case from a terrestrial campaign near Hanksville, Utah, USA and we validate a correlation panel generated with InCorr against a manually illustrated one, based on the same interpretation data. We further conducted a hands-on design validation to collect feedback from a broader range of geologists. In this way, we introduced interactive correlation panels to the domain of remote geological analysis, which is promising to foster traceability, reproducibility, and communicability of an otherwise static illustration.

## 4.2 Design Goals

Based on the previous observations we came to the hypothesis that the workflow of remote geological analysis would significantly benefit from an interactive and data-driven correlation panel, which achieves the following goals.

- **G3.1** The panel should be completely data-driven and consequently evolve with the analysis with minimal effort.
- **G3.2** The panel should allow geologists to relate visual representations in the correlation panel to the data they were created from.
- **G3.3** Such a panel should mimic manually illustrated correlation panels including stylistic freedom, but without overwhelming customization options.
- **G3.4** Using the panel needs to integrate well with the geologists' tool chain and workflow, otherwise it will not be used frequently.
- **G3.5** The panel should intuitively put 2D geological logs in context to each other and the 3D outcrop interpretation.

## 4.3 Related Work

The Petrel [Sch] software package for oil and gas exploration offers some outcrop measurement capabilities and correlation panels. As the software was originally intended for the analysis of seismic data in connection with drill shafts, the logs present in these panels are created from drilling wells and they have a different visual encoding than logs used for outcrop-based correlations. Therefore, they are not suitable for performing outcrop-based correlation analysis. General purpose GIS or 3D visualization tools, such as ArcGIS [Esr] or Cloud Compare [Clo], are commonly found in geological publications that include DOM analysis. These applications offer reliable measurement tools, but are neither targeted towards outcrop interpretation applications

nor do they support correlation analysis. Respective publications dealing with DOM interpretations typically describe a concatenation of data transformations [vLHRFP09, SG15]. A few specialized 3D outcrop interpretation tools have recently emerged, including LIME [BRN<sup>+</sup>19], VRGS [HGWR07], or PRo3D [VRVb, BGT<sup>+</sup>18]. All three feature a tool set for creating annotations and performing measurements on DOMs. Additionally, LIME and VRGS allow geologists to project manually illustrated logs onto the corresponding 3D surface. Nesbit et al. [NDH<sup>+</sup>18] use a 3D log in the context of their stratigraphic mapping. However, this logging is not integrated into an interactive workflow and the measurements need to be translated into 2D logs manually.

The depiction of geological phenomena has been an essential part of visualization research for decades, but mostly in the context of the analysis of seismic data for oil and mining. Patel et al. [PGT<sup>+</sup>08] introduce a tool to interpret 2D slices of seismic data from which they can pre-calculate horizon structures. Höllt et al. [HBG<sup>+</sup>11, HFG<sup>+</sup>12] present an interactive workflow for interpreting the 3D data directly by incorporating well logs retrieved from drilling. The steps of data retrieval, interpretation, well correlation and horizon extraction, and reservoir modeling [NLP<sup>+</sup>13] do align with the workflow of digital outcrop analysis, described in Section 4.4, however, the data, methods, and challenges differ significantly. Lidal et al. [LNP<sup>+</sup>13] focus on the communication of geological processes through visual stories and provide a sketch-based interface for the creation thereof. To the best of our knowledge, the field of digital outcrop analysis and the correlation of 3D outcrop interpretation data has not been explored or addressed by the visualization community.

## 4.4 Geological Analysis of Digital Outcrop Models

The basis of remote virtual outcrop geology is the acquisition of 3D data from rock outcrops. In this section we will discuss the development of outcrop observations to regional geological models. We focus on the stages relevant to our work and introduce the geological concepts that are important throughout this chapter. *Outcrop interpretation* (Section 4.4.1) produces contacts, which are the basis for *logging* (Section 4.4.2, Section 4.4.3). The result of the logging stage are geological logs, which are arranged in a correlation panel to discover and create *correlations* (Section 4.4.4). Correlations are used to reconstruct geological surfaces, which are then presented as 2D geological maps. We conclude this section with Table 4.1, summing up all relevant geological terms.

### 4.4.1 Outcrop Interpretation

Traditionally, geologists examine multiple outcrops in the field. At each outcrop they record textural characteristics, such as grain or crystal size and shape, relative color, layer thickness, layering patterns, layer orientation, and boundary geometries. Relevant measurements, photographs, notes, and outcrop sketches, as well as sketch logs or cross sections where necessary, are used to document the geology in the field, and analyzed out of the field for publication. The same process can be applied to DOM analysis. After measuring the height and width of an outcrop, geologists investigate *contacts*. A contact represents the delineation where one layer of rock, i.e., *stratum*, ends and another one begins. These contacts can be numerous and are

typically nested, so a stratum contains sub-strata, as shown by orange lines (*bedsets*) between red lines (*units*) in Figure 4.1b. Based on the different visible rock characteristics, geologists meticulously trace contacts by drawing polylines on the 3D surface along the discrete transitions between the deposited strata. This allows them to characterize the compositions of the strata of an outcrop hierarchically into sub-structures of arbitrary depth. Two contacts delimit a stratum that is homogeneous (to a certain extent). Geologists use different line thicknesses and colors to represent the magnitude of change between two adjacent strata. Geologists refer to this process of identifying contacts and hierarchical grouping as *geological interpretation*.

*Cross beds* are stratifications within a stratum, visible as thin white lines in Figure 4.1b. They are the preserved lee-faces of dunes or ripples, which migrated by wind or water action. The azimuth of the maximum dip direction of these cross beds indicates the original transport direction of the respective deposition medium [PS96]. The interpretation of cross beds is invaluable to determine the direction of wind or water flowing in the geological past. In the field, geologists use a compass-clinometer. In PRo3D, they trace a cross bed by picking 3D points on the DOM surface. Then a plane is fitted to these points via total-least square regression [JPJW16]. The result is a so-called *dip-and-strike* measurement, where the *dip* is the direction of maximum negative inclination of this plane, while the *strike* is orthogonal to the dipping direction, as illustrated in Figure 4.1c. In the context of correlation panels, geologists are primarily interested in the geographic direction of dips. Geographical directions are quantified as azimuth in degrees, where 0°, 90°, 180°, and 270° point to north, east, south, and west respectively.

The sizes of grains found in a stratum are a decisive factor, to determine its rock type and its mode of deposition, that is either by wind (aeolian) or by water (fluvial). Grain sizes range from coarse soil, such as cobble with 63-200 mm, to fine soil, such as clay with a grain size smaller than 0.002 mm. While in the field, geologists have many tools at their disposal, measuring grain sizes in DOMs is limited to grains visible on the textured mesh. For the exploration of Mars, scientists mostly have to rely on image-derived data, so they often need to infer grain sizes, rather than being able to measure them in 3D.

#### 4.4.2 Geological Logs

A 3D interpretation that characterizes an outcrop may contain a plethora of measurement values and annotations. For the sake of clarity, we will focus on the ones that are essential for creating a *geological log*, as depicted in Figure 4.2. A log characterizes the sequence of strata as they were deposited over time, starting with the oldest at the bottom and ending with the most recently deposited at the top. Perfectly horizontal strata are rare; geological effects and the roles of deposition and erosion in landscape evolution can produce irregular contacts. Therefore, the geologists ‘draw’ a log over the interpretation, connecting the contacts in their vertical, i.e., chronological sequence. The difference in elevation between two contacts determines the thickness of a stratum. Currently no 3D interpretation tool does support the direct semantic connection of contacts to create a geological log. Instead, elevation values and names of contacts are exported. Then geologists manually draw the corresponding log, while cross-checking with distance measurements in the 3D visualization. The notion of *true thickness* complicates this matter significantly, which we discuss in Section 4.4.3.

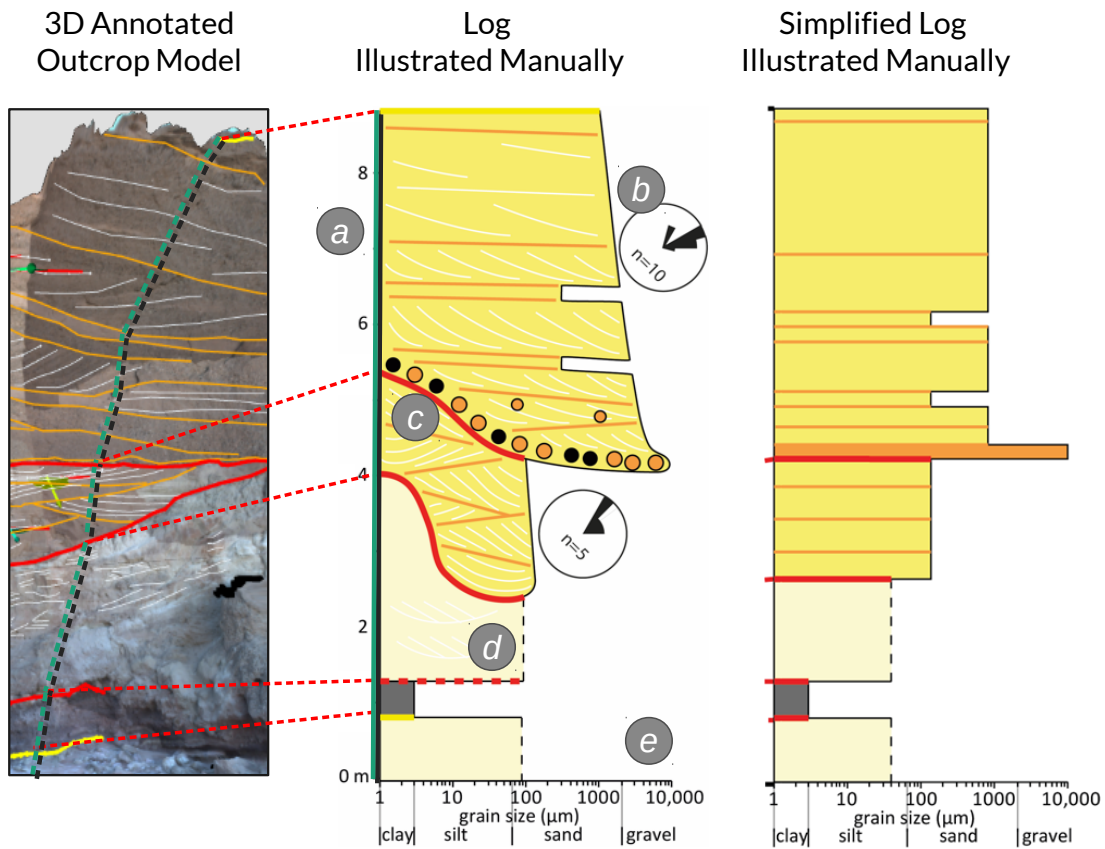


Figure 4.2: Annotated 3D outcrop model, followed by a sophisticated and a simplified log representation, both illustrated manually with a vector drawing tool. (a) Vertical axis encodes true thickness of strata. (b) Rose diagrams show the distribution of dipping orientations. (c) Individual styles and glyphs convey rock characteristics. (d) Dashed lines convey uncertainty. (e) Horizontal axis encodes grain sizes logarithmically, directly relating to rock types.

A geological log characterizes an outcrop by showing the type and succession of strata in an abstracted form. As illustrated by Figure 4.2a, the vertical axis of the log encodes the elevation of the contacts measured in the 3D view as well as the thickness of the strata enclosed by the respective contacts. The horizontal axis in the log (Figure 4.2e) corresponds to the grain size on a logarithmic scale. Grain size relates to the rock type of a stratum, which is also encoded in the stratum's color. To characterize the orientation of a stratum, geologists visualize the distribution of cross bed dipping-azimuths in a rose diagram. In the example in Figure 4.2b, the upper unit between the yellow and the red contact 'dips towards east/north/east' based on ten dip-and-strike measurements. Dashed lines are often used to convey uncertainty, for instance, concerning a contact (horizontal line) or concerning the grain-size (vertical line) shown in Figure 4.2d. Geologists use additional encodings such as curved lines for contacts of varying elevation or glyphs to convey grain distributions (Figure 4.2c).




Figure 4.3: The apparent thickness of a stratum observed at an outcrop is often misleading. The rock may be broken off at a slanted angle or the stratum itself may be tilted. A dip-and-strike measurement is essential to determine a stratum's orientation and compute its true thickness.

#### 4.4.3 True Thickness

For the sake of simplicity, we accepted that the difference in geographic elevation between two contacts results in the thickness of the enclosed stratum. In nature that is often not the case, which is why geologists distinguish the measured or *apparent thickness* and the *true thickness* of a stratum. Our previous simplification is only valid if the measured strata lie in a horizontal plane and the outcrop surface is vertical. Figure 4.3 illustrates the discrepancies between geographic elevation, apparent thickness, and true thickness. In reality, the deposition of material originally occurs horizontally, but strata may be tilted or even folded over by a variety of geological or geomorphological phenomena. Hence, instead of using global elevation values over all strata and logs, each apparent thickness needs to be corrected by the stratum's dipping angle. In stratigraphy it is an accepted simplification to use one angle per log, which still results in each log creating its own coordinate system. Throughout this chapter, when we speak of thickness, we mean true thickness, since the apparent thickness is of little relevance for geological interpretation.






Figure 4.4: Logs arranged in a manually illustrated correlation panel with (a) converging contacts, (b) spatial distances between logs along geographic direction, and (c) a fossil layer used as leveling horizon.

#### 4.4.4 Correlation Panel

After having created multiple logs, geologists arrange them in a correlation panel, as shown in Figure 4.4. The juxtaposition of logs allows geologists to identify similar strata across outcrops and to *correlate* them, hence the name correlation panel. The notion of geological correlation is not related to the mathematical concept. Correlations are visualized as colored connections between contacts. A dashed pattern is used to express uncertainty if two contacts belong to the same structure or not. The arrangement of logs from left to right is often determined by their succession along the course of a geomorphological feature, such as a canyon or a crater rim. Rulers between the individual logs encode the distances between them (Figure 4.4b). The distances also convey a degree of uncertainty and data quality, since inferring correlations between outcrops across large distances is less reliable. Due to geological faults or other geomorphological processes, different outcrops do not necessarily expose the same stratum at the same elevation or with the same thickness. Consequently, correlation lines are rarely horizontal, which is why geologists often choose a distinct contact as a leveling horizon for all outcrops. In the example shown in Figure 4.4, a stratum rich of fossils has been found in the first and the last outcrop, indicated by dark gray rectangles at an elevation of 1m. This stratum is missing from the other logs, because it is not exposed at these locations and potentially buried. Still, geologists are able to infer its existence and its approximate elevation.

#### 4.4.5 Hanksville-Burpee Dinosaur Quarry, UT, USA

Image data was collected at the Hanksville-Burpee Dinosaur Quarry (HBDQ), near Hanksville, Utah, U.S.A. (110°47'30"W, 38°27'12"N). The distinctive rocks of the Jurassic Morrison Formation have a sub-horizontal regional dip, and are therefore ideal for testing the capabilities of PRo3D in manually creating and correlating geological logs. During the campaign, the geologists digitally captured four outcrops along a canyon, as illustrated by the map in Figure 4.1a, named Sol2–5. After capturing, they reconstructed DOMs using PRoViP [Gao16] and interpreted each of them in PRo3D yielding contacts and cross-bed measurements. Based on the interpretations they manually measured the true thicknesses of the strata and drew a log for each outcrop. All outcrop interpretations of this dataset do contain two specific strata, *units* enclosed by unit contacts in red and *bedsets* enclosed by bed contacts in orange, as illustrated in Figure 4.1. They arranged the individual logs in a manually illustrated panel and preliminary derived correlations. The resulting correlation panel is visible in Figure 4.4. We use the DOMs, the interpretations, the logs, and the correlation panel as a running example throughout this chapter and as a use case in Section 4.7.

### 4.5 Design Process and Domain Abstraction

The results of InCorr are based on a decade-long collaboration with planetary geologists, mainly in the context of research and development of PRo3D as an interpretation tool. In a workshop following the evaluation campaign described in Section 4.4.5, our collaborators stated the need to semi-automatically generate correlation panels from interpretations and suggested a simplified visual encoding for logs, as shown in Figure 4.2. We followed a participatory design approach [JKKS20] leading to a three-phase evolution of InCorr: In Phase (1), we were

outcrop	exposed rock face showing sedimentary structures
stratum / -a	layer of rock bounded by two contacts
contact	discrete transition between two strata
dip	inclination angle of a stratum
dip-and-strike	measurement to determine the orientation of a stratum
true thickness	thickness of a stratum with respect to its dip
cross bed	cross stratifications within a stratum
unit, bedset	specific strata, where a bedset is a sub-stratum of a unit

Table 4.1: Summary of the most important geological terms.

concerned with how to transform annotations into a log, researching a hierarchical data structure and the transformations necessary. This resulted in a non-interactive log prototype matching the simplified visual encoding. In Phase (2), we focused on understanding the domain background and tasks involved with correlation analysis and created an interactive prototype. We integrated it with PRo3D featuring multiple logs and correlations. In Phase (3), the necessity of a logging tool measuring true thickness became evident. Key aspects of this phase were interaction and visualization design, and end-to-end workflow integration. Each phase was accompanied by a week-long research stay and roughly quarterly meetings. During Phase (3) we shortened intervals and iterated on visualization and interaction prototypes sometimes on a daily basis, shaping InCorr through the continuous feedback provided by our collaborators.

#### 4.5.1 Task Analysis

With InCorr we address a set of tasks that bridges the gap between outcrop interpretation and the creation of a geological model based on logs and correlations:

- **T1** Create a geological log for each outcrop based on the annotations and measurements taken.
- **T2** Create correlations from geological logs as the basis for a regional geological model.
- **T3** Edit and export the correlation panel to be manipulated in other tools for further analysis or dissemination.

We abstract these tasks and their subtasks by following the multi-level task typology by Brehmer and Munzner [BM13]. We subdivide the creation of a geological log (**T1**) into connect contacts to create strata (**T1a**), assign rock types to strata (**T1b**), and assign cross beds to strata (**T1c**). In **T1a** the geologists *identify* the relevant contacts they want to connect and then *select* them to form a log (*annotate*). For each stratum in the log they *identify* the grain size

and select the respective rock type (**T1b**). Further, they identify and select cross bed measurements belonging to a stratum to summarize the orientation distribution within it (**T1c**). For **T2** the geologists arrange the logs from **T1** to compare their characteristics and find similar strata (**T2a**). When found, they select one contact per log, and connect contacts to create (annotate) correlations (**T2b**). After the correlation analysis is completed, geologists ‘tidy up’ the correlation panel. They vertically and horizontally arrange logs and export the correlation panel to be used in other tools of their workflow (**T3**).

In relation to multivariate geometric tasks (Section 1.1.2), **T2 Identification** and **T5 Annotation** are predominant in this setting. Geologists inspect high-resolution DOMs to identify geological features that support the selection of contact delineations, rock types, and cross-bed measurements. The manual creation of the underlying interpretation data heavily relies on **T2** and **T5** as well.

### 4.5.2 Data Transformation

To bridge the gap between outcrop interpretation and 3D logging, we need to infer the hierarchical structure of strata from a set of contacts. Logging allows geologists to pick a 3D position on each contact, while the true height of these positions, i.e., elevation corrected by a dipping angle, determines the vertical sequence of the contacts. Each contact has a rank assigned by the geologist, representing the magnitude of change between two rock layers. With this information, we can derive a tree of strata. Each stratum is bounded by an upper and lower contact defining its minimum and maximum height, and thus its thickness. We achieve this by starting out with a fictive stratum with a height interval from negative to positive infinity. Our algorithm goes through the set of contacts sorted by their rank, finds the stratum with a height range matching the contact’s height and splits it into two sub-strata. Performing this step for each contact yields a tree of strata in a breadth-first fashion.

Each stratum can be assigned a rock type and may contain geological features, such as cross beds. To characterize the distribution of cross bed orientations within a stratum, geologists perform dip-and-strike measurements. Such measurements can then be assigned to leaf strata. This allows us to use simple tree traversals to aggregate measurements for arbitrary strata. As an example, we represent the log illustration shown in Figure 4.2 using our data model in Figure 4.5a. The whole outcrop, represented as root stratum, contains units (sub-strata) defined by the red contacts. The upper two units contain sets defined by orange contacts, which do not have descendant strata, but contain cross beds (white lines). Dipping-azimuth values of the cross beds are aggregated at a unit level and visualized as rose diagrams. We define a correlation as a set of contacts connecting them between logs. Correlations are often uncertain in parts, which means that the connection between two contacts is either certain or uncertain, typically represented by a solid or a dashed line, respectively.

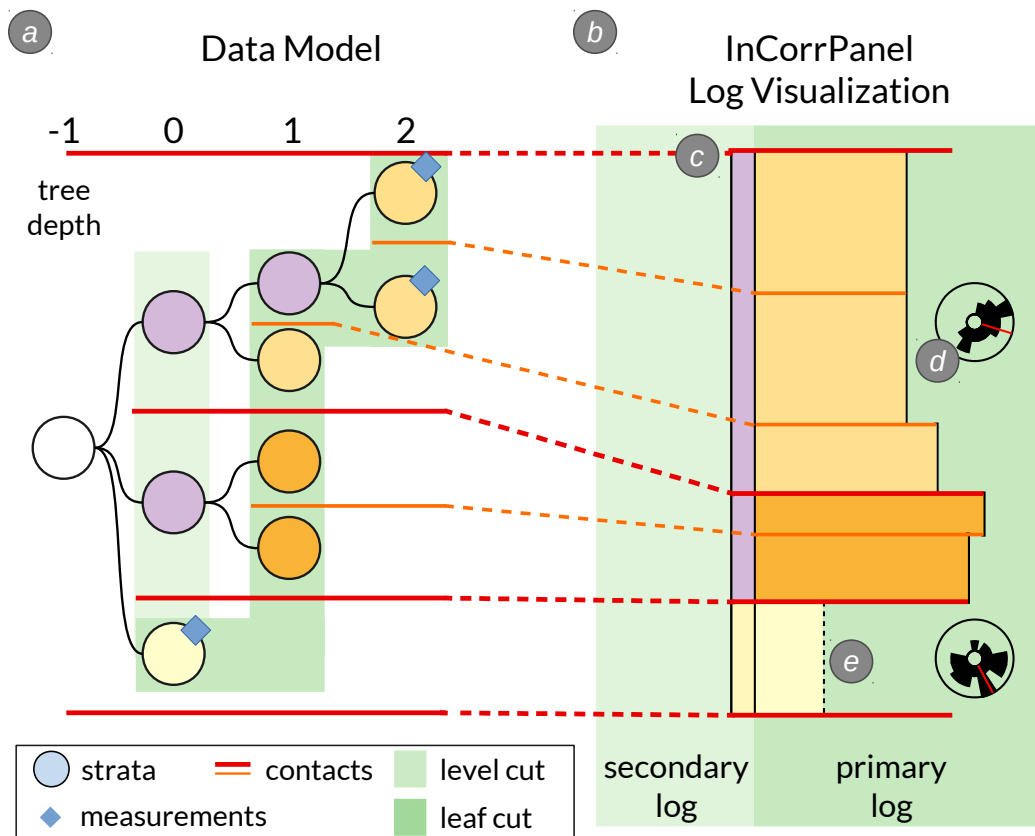


Figure 4.5: (a) Data model: tree of strata and their bounding contacts. (b) InCorrPanel log: (c) secondary log shows a tree cut at depth 1, while the primary log shows the leaf cut. (d) Rose diagrams aggregate cross bed orientations. (e) Dashed border conveys uncertain an rock type.

## 4.6 Visualization and Interaction Design

InCorr consists of three components that we integrated with PRo3D, shown in Figure 4.6. First, the *InCorrPanel*, a 2D interactive correlation panel, that offers geologists an evolving summary of their geological analysis and that is easy to keep in sync with annotations and measurements, second, a *logging tool* that allows geologists to intuitively connect contacts in the Outcrop View provided by PRo3D (see Figure 4.6a). Third, a *list view* for assigning rock types to strata in the InCorrPanel. In the following, we discuss the visual encodings in Section 4.6.1 and the interaction design in that we use for these components Section 4.6.2.

### 4.6.1 Visual Encodings

Geological logs are a visual abstraction of complex spatial relationships of 3D phenomena. Showing them in a correlation panel does provide geologists with a concise overview of large-scale geological analyses. Static correlation panels already facilitate **T2** by enabling a visual

#### 4. INCORR: INTERACTIVE DATA-DRIVEN CORRELATION PANELS FOR DIGITAL OUTCROP ANALYSIS




Figure 4.6: System overview of InCorr: (a) Outcrop View, showing a digital outcrop model interpreted in 3D by a geologist and annotated with the vertical 3D logging tool indicating layer thicknesses. (b) InCorrPanel, showing logs for two outcrops created directly from annotation data mimicking manual illustrations. (c) GUI to assign rock types to rock layers.

comparison of outcrops without inspecting their spatial representations. To preserve the benefits of static illustrations, the design of the InCorrPanel needs to closely resemble manually illustrated logs and correlation panels, without overwhelming users with customization options (**G3.3**).

There is no universally agreed upon standard for creating geological logs. Every geologist has their own specific style. For instance, the geological logs in our running example (shown in Figure 4.4) can be considered a rather sophisticated variant. We did not conduct a formal design space analysis, but let our collaborators decide on the appropriate visual representations. A key topic of design phase (2) was to derive which aspects need standardization and which aspects need to be left to the individual geologist. As currently there is no tool for creating correlation panels directly from outcrop interpretations, our collaborators were eager to participate in design sessions. Hayes et al. [HGE<sup>+</sup>11], Van Lanen et al. [vLHRFP09], or Hampson et al. [HGS<sup>+</sup>11] present examples of the variety of correlation panels published in geological journals. In the following we present our design decisions for InCorr starting with the visual encodings to display a single log.

### Single Log

The visual representation of a single log in the InCorrPanel is illustrated in Figure 4.5b. The horizontal axis encodes the grain size on a logarithmic scale, while the vertical axis encodes the order and true thickness between contacts. To calculate the true thickness, a log requires a reference plane in 3D. For each contact we compute the height above the reference plane and map it directly to a position on the vertical axis. For each stratum, we draw a filled rectangle, ranging on the vertical axis from the elevation of its lower contact to the elevation of its upper contact. The stratum's grain size, which relates to the associated rock type, is encoded redundantly by the width of the rectangle and its color. Additionally, we draw the contacts as horizontal lines at their elevation matching their thickness and color in the Outcrop View (**G3.2**). To indicate if geologists are uncertain with the rock type they assigned, we draw the right border of the rectangle as a dashed line, as in Figure 4.5e. This is especially relevant in Martian use cases, where the tools for observation are mostly visual and the better part of the observed strata are of finer grains. We chose the simplified rectangular log over the curved contact representations and curved polygons for several reasons. One could derive the variation in elevation by picking multiple points per log, but (1) there is no information in a contact that would describe the variation of a stratum along the horizontal axis; (2) it becomes difficult to read strata thicknesses directly from the log; and (3) the encoding is rather specific and rarely used in manual illustrations. Ultimately, as we create the plot as an SVG image, it can be directly edited in a vector graphics tool after export to add custom illustration styles (**G3.3**, **G3.4**).

Embedding an outcrop analysis into a bigger context, deep hierarchies of strata may emerge. To manage such hierarchies, our collaborators suggested to add a second, abstracted log representation that emphasizes the affiliation of sub-strata to their respective containing strata. We denote this abstracted log as *secondary* log and show it to the left of the *primary* log (Figure 4.5c). The primary log always shows the full detail of the strata (leaf nodes of the data model), whereas the secondary log shows a single, user specified level of the hierarchy indicated in the data model (level 0 in Figure 4.5a). This two-column representation of logs resembles a specialized icicle

plot [Sch11], where all nodes except the leaf nodes and the selected level are removed from the hierarchy. A similar double log representation can, for instance, be found in the manually illustrated correlation panel by Hampson et al. [HGS<sup>+</sup>11] (**G3.3**).

The selected depth in the secondary log also governs the granularity of how we aggregate measurement values. Supporting **T2**, the InCorrPanel summarizes the distribution of dip-azimuths as rose diagrams (Figure 4.5d). We aggregate the orientations into 24 15° angular bins and encode the frequency into their area, as suggested by Sanderson and Peacock [SP19]. This allows the geologists to infer the major dipping-azimuths of a stratum without missing low frequency outliers. We further compute the mean angle (via polar coordinates) and encode it as a red line. According to the geologists, this representation allows them to quickly judge if the distribution follows one general direction.

### Multiple Logs and Correlations

In published correlation panels, the logs are carefully ordered to convey the results in the best possible way. In the case of the HBDQ analysis, the logs are sorted along the geographic direction from SSW to NNE, as the result of a projection from longitude and latitude onto a line, as illustrated in Figure 4.1a. We did consider ordering the logs in the InCorrPanel automatically along a fitted line or letting the geologists pick a geographic direction. Ultimately, the log order is very context sensitive to the location of the outcrops and domain experts will want to order them manually. Each log is a concise summary of an outcrop interpretation, and a correlation panel supports geologists in identifying strata with similar characteristics. The InCorrPanel allows them to juxtapose logs arbitrarily in an interactive fashion. This enables them to quickly compare rock types and cross bed orientations during the analysis in order to correlate contacts (**T2**). This is similar to rearranging the dimensions of a parallel coordinates plot to investigate relationships. Furthermore, we compute the spatial distances between adjacent logs and display them as rulers between the log names.

After determining the horizontal arrangement, we are left with the vertical positioning of each log in the InCorrPanel. In discussion with our collaborators, we realized that there is no single layout method for vertically positioning logs. Additionally, the most suitable layout is likely to change throughout the analysis. We agreed on a two-fold approach, which provides standardization as well as flexibility (**G3.1**, **G3.3**). We started out rendering all logs, their contacts, and strata in a common coordinate system based on the geographic elevation of the individual elements. Since this approach does not accommodate for true thickness we use the geographic elevation of a log's reference plane as an anchor point as discussed in this section. When users find correlating strata they often use the resulting surface as baseline for vertical alignment. In the best case, every log, i.e., every outcrop, exhibits a part of the same rock stratum. Using this stratum as the baseline offsets each log vertically in such a way that all log nodes of this surface are at the same vertical position in the log. In that case, the correlation lines associated with the upper contact of the stratum become horizontal. In the example in Figure 4.4c, only two outcrops exhibit the baseline stratum. The dashed yellow lines in the correlation panel indicate that the correlated surface is also present at the locations of the other two outcrops, but not exposed.






Figure 4.7: To create a 3D log, users (a) first pick a reference plane, (b) then they connect the contacts to form (c) a 3D log resulting in an initially (d) empty log.

#### 4.6.2 Interaction

In this section we discuss the interactions necessary to achieve the domain tasks in the sequence described in Section 4.5.1, i.e., creating logs (**T1**), creating correlations (**T2**), and editing the resulting panel for export to other tools (**T3**). Figure 4.7 and Figure 4.8 provide an overview of the succession of interactions for creating a log (**T1**), starting with the creation of a single log based on 3D contact lines from an outcrop interpretation (**T1a**).

**T1a Connect Contacts to Create Strata.** Before users can start connecting contacts to form a log, they need to specify a reference plane to enable true thickness computation. They choose a contact which they assume to have a suitable orientation to which a plane is fitted using 3D total-least squares regression. The fitted plane is then visualized as a colored disc, where the color encodes the dipping angle of the plane (Figure 4.7a). Geologists can try out different contacts and inspect the respective planes or confirm the found plane. Then users pick a point on each contact they want to add to their log. Due to the fact, that contacts are implicitly ordered by




Figure 4.8: Users assign rock types to the empty log via the (a) rock types list, resulting in (b) a log encoding rock categories and grain sizes. To complete the log, geologists (c) assign cross bed measurements, which are aggregated and encoded as (d) rose diagrams.

their elevation, the selection of contacts needs not to happen in order. For each picked point we compute its height above the reference plane and sort the set of points by height over the plane. This allows us to draw 3D line segments by creating pairs of points adjacent in the sorted list (Figure 4.7c). Further, a point picked for a contact can be modified or removed and the log polyline changes accordingly, giving immediate visual feedback to the user (Figure 4.7b). Users can edit existing logs in the same way to keep their analysis up-to-date with new contacts (G3.1).

**T1b Assign Rock Types to Strata.** The strata of a newly created log do not have rock types associated with them yet, which is indicated by their white color (Figure 4.7d). To assign a rock type to a stratum, users first select the respective stratum in the primary log and then they pick a rock type from the rock types list (Figure 4.8a). The color and the width of the respective rectangle changes accordingly. To indicate how sure geologists are with the choice of the rock type they can add an uncertainty indicator as a dashed or solid right border encoding uncertainty or certainty, respectively. A log completely annotated with rock types, including uncertainty indicators, is shown in Figure 4.8b (G3.3). Selecting a stratum in the InCorrPanel, we also

highlight the bordering contacts in the Outcrop View (Figure 4.8c), allowing geologists to identify the respective area of the outcrop (**G3.5**).

**T1c Assign Cross Beds to Strata.** Analogously to the previous task, users first select a stratum in the primary log of the InCorrPanel (Figure 4.8d) and the bounding contacts are highlighted in the Outcrop View. This supports the geologists in identifying the region where they want to select cross bed measurements (**G3.5**). After confirmation, the selected cross beds are assigned to the selected stratum and a rose diagram appears next to the log. From this point on, selecting the stratum also selects the cross beds assigned to it. Changing the set of selected cross beds by adding or removing items is also reflected in the rose diagram (**G3.1**). We discussed approaches for the automatic selection of cross beds with our collaborators, but in the end resorted to letting them select the measurements manually. After having created at least two logs, i.e. performing **T1a** twice, users can create correlations (**T2**).

**T2a Find Similar Strata.** Finding correlations means that geologists need to identify contacts that belong together, i.e. two contacts delineate the transition between the same strata in different logs. To achieve this, geologists need to be able to effectively compare outcrop characteristics encoded in a log, comprising the rock type of a stratum, its thickness, and the orientation distribution of the cross beds within it. To support this task, the InCorrPanel allows users to arbitrarily change the horizontal order of logs and compare any two logs via juxtaposition.

**T2b Connect Contacts to Create Correlations.** When corresponding contacts are identified, the geologists select the respective lines, one per log. On confirmation, the selected contacts are then connected by curved lines in the color of the selected contacts. Each connection between two correlated logs is initially rendered in a dashed pattern to convey uncertainty and can be switched to a solid line style. If the correlation analysis is complete and the InCorrPanel contains all the findings the geologists want to present their results for scientific dissemination, for instance as an essential part of a geological publication, they move on to **T3**.

**T3 Arrange Logs and Export the InCorrPanel.** To make logs more comparable in a published correlation panel, geologists typically level all logs to a common horizon (**G3.3**). This means to vertically arrange and align all logs to a common baseline correlation. When drawn in the same coordinate system, correlating contacts rarely occur on the same height, as exhibited by the red correlations in Figure 4.4. Ideally there is a distinctive rock layer present in every log, which results in a correlation across the whole panel. Then this correlation can act as a vertical baseline. To achieve this, geologists select a correlation in the InCorrPanel and confirm to vertically align all logs to it. This results in straight connections instead of curved ones, as illustrated by Figure 4.10e. For final adaptations and editing, the content of the InCorrPanel can be exported as a scalable vector graphic readable by all common vector drawing applications (**G3.4**).

### 4.6.3 Implementation

InCorr is implemented in F# as an extension of PRo3D [VRVb], utilizing the ELM-style framework Aardvark [VRVa] for building scalable visual computing applications. Programming types and transformation functions follow the rules of domain driven design [Wla18] and align with the data model described in Section 4.5.2. The InCorrPanel is created in the SVG format, allowing

geologists to easily refine the exported visualization in a vector illustration software such as Adobe Illustrator. The 3D view, GUI elements, and the InCorrPanel are composed via a JavaScript code generator to run in a web browser.

## 4.7 Evaluation

In this section we discuss the methods we used to evaluate InCorr. We verified InCorr by using it to recreate the correlation analysis of the HBDQ campaign. Then we validated the resulting InCorrPanel (Figure 4.10a) against the manually created illustrations shown in Figure 4.10b–d. Finally, we collected feedback from three geologists after performing tasks in the form of a design validation.

### 4.7.1 Use Case - Hanksville-Burpee Dinosaur Quarry

We used InCorr to create 4 logs, one for each outcrop of the HBDQ dataset named Sol 2–5, as shown in Figure 4.9. We used the logging tool to connect 62 contacts from the interpretation data, which amounted to 58 true thickness measurements. We assigned the rock types to the same number of strata and selected 18 strata to which we assigned a total of 129 cross bed measurements. We then added 2 correlations, which resulted in the correlation panel presented in Figure 4.10a. The whole process took approximately 1.5 hours. Carrying out this process in a non-assisted manner, using field data, involves manual measurement of the distances between contacts, conversion of apparent to true thickness, and manual plotting of rose diagrams. According to our collaborators, for the amount of data presented, this would take considerably longer than 1.5 hours, and though the time it would take would vary between workers, producing a correlation of equal precision would be on the order of a day to several days of work.

Our log representation matches the visual properties of the simplified manually illustrated log (Figure 4.10d) and so meets the minimal encoding capabilities, as proposed by our collaborators (Section 4.6.1). We extended this version by two data-driven features, the secondary log and the rose diagrams. Our solution does not include the following styles present in the manually illustrated log (Figure 4.10c): (1) curved, converging, or tilted contacts, (2) gradients, pattern fills, and rounded corners for strata, nor (3) glyphs to indicate additional rock properties. Besides depending on the individual geologist’s preferences, most of these features do not directly encode information that is readily quantifiable in the data. We left these encodings to the manual workflow following the export of the InCorrPanel. Nevertheless, it is interesting to incorporate styles that could be inferred from the data, such as curved or tilted contacts. We could not recreate the thicknesses of one stratum in Sol 2 and one in Sol 5 (Figure 4.10f). The reason is either vertical exaggeration to make them visible or the true thickness is inferred from information not present in the data. When we compare the manually illustrated panel (Figure 4.10b) with the InCorrPanel as a whole, the main difference is that our illustration of the HBDQ campaign is flattened to the top correlation instead of the one connecting the two fossil layers.




Figure 4.9: Outcrops Sol 2–5 of the HBDQ dataset, each with a 3D log created with InCorr and the resulting 2D log from the InCorrPanel.

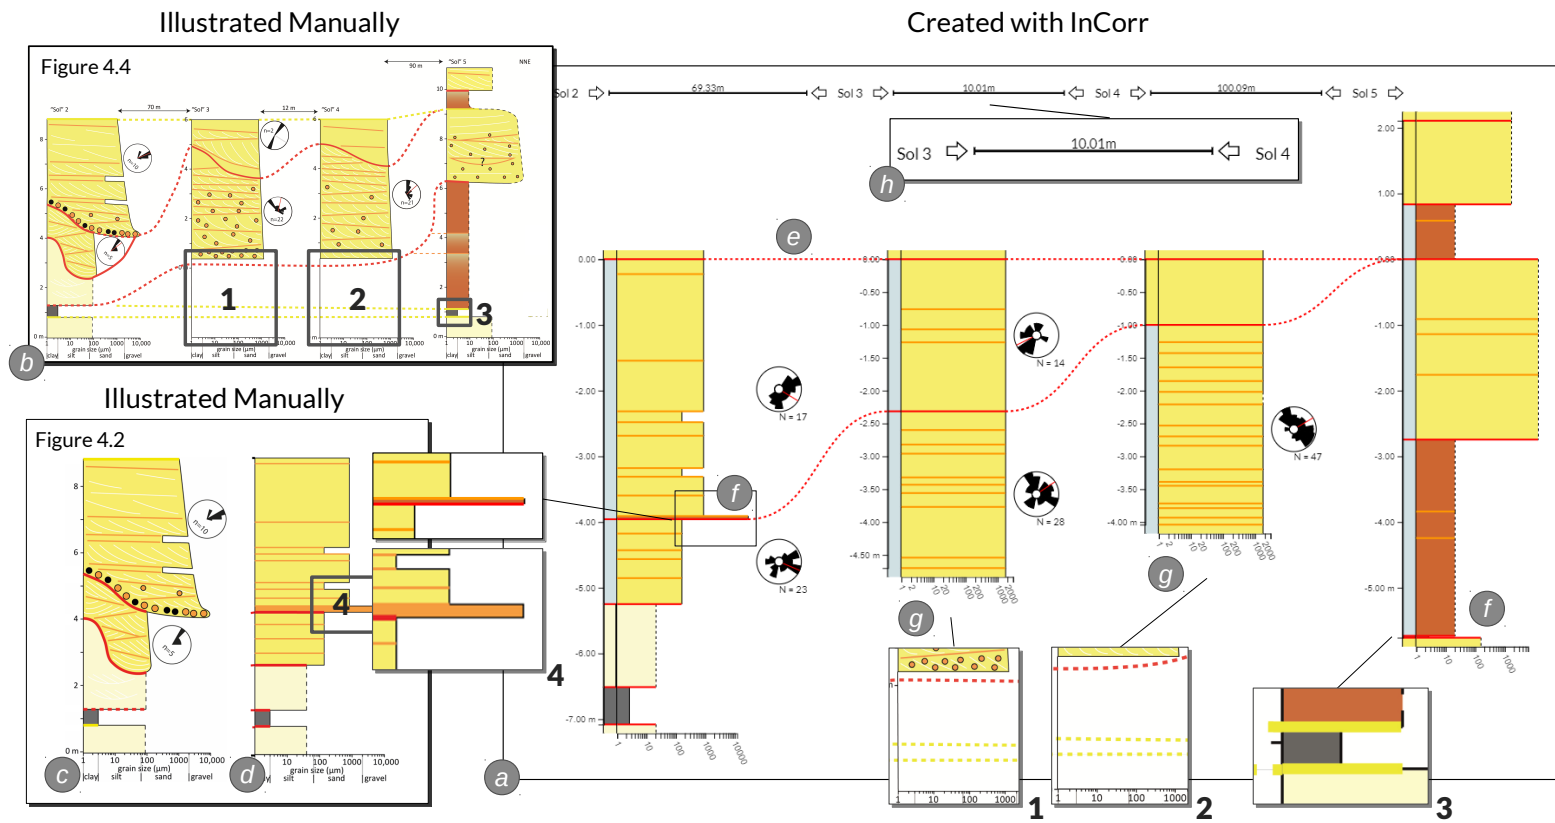


Figure 4.10: (a) Correlation panel created with InCorr based on outcrop interpretation data with (b,c,d) manual illustrations with detail lenses (1-4) for comparison. (e) All logs are aligned to a common baseline correlation at the top. (f) No vertical exaggeration of strata, (g) no indication of buried regions, and no visible connection of non-adjacent logs. (h) Rulers showing spatial distance between logs.

This is because in the manually illustrated panel the location of the fossil layer is assumed for Sol 3 and Sol 4 although it is actually buried in the physical outcrop and therefore the logs contain an empty region (Figure 4.10g). We did not include support for specifying hidden strata in InCorr, because it would require additional interactions and design iterations, and geologists can rarely make such assumptions on Mars. We did not implement the visual connection of contacts with a line between non-adjacent logs. This would require a layout algorithm that avoids collisions with logs in-between. At the same time, the resulting curved line should encode the course of the correlation, which can only be specified by the geologist. But, a correlation can be created between arbitrary logs, and the visual connection will show when the respective logs are adjacent.

### 4.7.2 Design Validation

We recruited three geologists to participate in our design validation and we will refer to them as P1, P2, and P3. P1 is a sedimentologist involved in design phases (1) and (2), P2 is an engineering geologist from the field of tunnel constructions, and P3 is a planetary scientist with a geological background, none of which were co-authors of the paper this chapter is based on. We conducted one-to-one interviews limited to one hour, starting with background questions about their field and their expertise in DOM interpretation and logging. Then we asked them to perform the tasks of creating logs **T1** and creating correlations **T2** on two outcrops of the HBDQ dataset, closing with a short questionnaire and open feedback. Prior to the interview the participants received a five minutes tutorial video, briefly explaining the interactions for achieving **T1** and **T2**. We decided to collect qualitative feedback on the relevance of the problem, the quality and ease of our solution, and its potential impact.

All participants agreed, that logging and the correlation analysis of DOMs is a time-consuming process: *‘Often it is faster to draw the panel by hand on scales paper (opposed to using vector drawing software)’* (P3), and that InCorr can provide a significant speed up to logging and assigning cross beds: *‘Looking at this software I should be able to do it very quickly. About an hour instead of a week.’* (P1).

They were very positive about the interactions in the 3D Outcrop View while creating a log and easily changing the succession of points: *‘Usually it is hard to get the scales correct’* (P3), and *‘That works pretty well!’* (P1). The linked highlighting was also a very welcome feature for relating logs and their spatial counterparts, but also when assigning cross beds: *‘The highlighting of a stratum’s contacts is very intuitive. It feels visual and haptic.’*, *‘This is perfect for correlating spatially over large distances.’* (P2).

Each participant was missing different stylized features in the logs and correlations representation. Asked if they could encounter the InCorrPanel ‘as is’ in a geological publication they generally agreed, but they themselves would not use it directly without additional editing. Although P1 was content with post production: *‘Really good that it is already in SVG, that’s important. Glyphs and styles don’t need to happen in InCorr, I would export it to AI or Corel anyway’*. Still, he wished for more options on encoding different properties along the correlation lines. If two correlation lines end in one, it is interesting to show where the phase out happens. P2 wished for hatching patterns to fill strata and to be able to change the roundness of the rectangles to encode weathering

effects. He also pointed out that he is unfamiliar with the quantification of the grain size in our rock types list. With our collaborators we agreed on displaying grain size categories as values of the Krumbein  $\phi$  scale [KS64], but for broader acceptance we should also show the metric values and support other categorizations. P3 immediately addressed the absence of glyphs and curved contacts.

When asked about the potential and implications of the tools provided through InCorr, each participant expressed their excitement, but for different reasons. P1 plans to use the logging tool for a publication requiring an extensive true-thickness analysis. P2 is eager to try out logging and correlation analysis in lithological outcrops, i.e. tunnel faces created during tunnel excavation. P3 focused on the implications of a data-driven approach to correlation analysis, since in her current research she needs to reproduce and compare results from data that is only present in figures. On the suggestion of bundling the analysis and input data with a publication to be viewed with InCorr and PRo3D, she answered *'Isn't this the whole point of a publication? To have other people interact with your data. I hope that is the future of publishing'*.

### 4.8 Summary

In this section we summarize the key aspects of InCorr, and relate them to the design goals we defined in Section 4.2. We present a critical appraisal of our solution supported by the results from Section 4.7, referring to participants of the design validation where necessary.

We developed a 3D logging tool that allows geologists to create geological logs with little effort, especially when compared to current approaches (**G3.1**). All participants could create a 3D log in under a minute, while spending most of the time on considering which contacts to include. We had some usability issues when selecting contact lines or very thin strata in the InCorrPanel. To improve the selection process of thin lines, a list of potential candidates could be provided, combined with additional visual feedback, like a selection preview while hovering the lines or thin strata. Assigning cross beds to strata would greatly benefit from better selection tools provided by PRo3D, as for instance a line brush. Further, with the knowledge gained throughout this work we could try again to offer our collaborators an automatic pre-assignment of cross beds within a stratum that they can then intuitively edit.

Our data-driven approach of logging directly on the 3D contact lines ties logs and correlations to the outcrop interpretation data (**G3.2**). All participants recognized this as a significant improvement over current methods, where a log is rather a collection of height measurements than a 3D polyline. It would be interesting to investigate how to enrich the exported vector graphics file with meta data indicating its origin. According to the design feedback, the InCorrPanel contains a sufficient set of visual encodings and customization to pass as a correlation panel in a geological publication (**G3.3**). In Section 4.7 we could demonstrate, that InCorr encompasses all features to conduct a correlation analysis. Nevertheless, **G3.3** offers the largest room for improvement. The InCorrPanel establishes a baseline, while some styling features are clearly left to vector drawing tools. It would be beneficial to explore the boundary of this separation, adding features that are supporting the analysis or should be reflected in the data.



The integration with PRo3D and the export of the InCorrPanel to a vector format allows InCorr to successfully bridge the gap between outcrop interpretation and creating regional geological models based on correlations (**G3.4**). We also verified this by reconstructing the correlation analysis of the HBDQ from existing outcrop interpretation data. On the other end, Figure 4.5 shows a modified log representation exported from the InCorrPanel. The highlighting of the selected log and of a stratum's bordering contacts allows users to establish a context between the 3D outcrop interpretation and the 2D log representation (**G3.5**). Geologists are very aware of the spatial context between the outcrops and logs of their analysis, due to the fact that they may spend weeks on interpreting the outcrops. None of the participants was familiar with the HBDQ data, so the contact highlights helped them to quickly orient themselves in the interpretation. We noticed that P3 subsequently tried to select strata in the 3D Outcrop View, i.e. the region on the surface between two outcrops. Geometry for strata, which one could infer from contacts and the log, would benefit InCorr on multiple levels. It would be the basis for the assisted assignment of cross beds as mentioned before. One could directly assign rock types in the 3D Outcrop View without switching to the panel, and further it could tie the actual DOM geometry and texture to the logging. P1 also entertained the idea of juxtaposing outcrops in 3D, aligning them along a geographic direction, analogous to the correlation panel. Such a transformation would dissolve longitude and latitude components of the data, while preserving elevation and orientation.

InCorr was designed to support geologists in the geological analysis of digital outcrop models by translating static correlation panels to the interactive world. We could identify several other application areas for InCorr. As suggested by P2, InCorr can be used to log the succession of strata in solid rock, exposed in the form of tunnel face outcrops. These are created and captured at regular intervals during tunnel excavation. In large-scale tunnel construction projects it is common practice to maintain a reconnaissance tunnel, smaller and a few hundred meters ahead of the main tunnel. Logging and correlation of the reconnaissance tunnel allows geologists to reconstruct the geological situation in the mountain and predict when the main tunnel will hit critical rock sections. Outcrop-based logging and correlation panels are not used in this field of geology and we will pursue this idea in an ongoing research project. In the context of sedimentary geology and the exploration of Mars, we can adapt the InCorrPanel to display other measurement aggregations of strata than dip azimuths. Our data model and interactions support to add, for instance, bedding thicknesses or grain sizes to strata, and show their distribution in histograms. Finally, stratigraphic succession plays an essential role in archeological excavations and InCorr could be a valuable addition to the work of Traxler et al. [TN08] on archeological stratigraphy.

Although correlation panels are a very domain-specific representation, we can generalize them as follows: a log connects discrete features to form intervals. The order is implicitly given by a metric and the intervals are nested through the rank of their features. Users connect two features in different logs through a correlation depending on the similarity of the two intervals the feature is part of. As an example, one could use the InCorrPanel to compare the lives of different people. Each person's life consists of events with different magnitudes of change. These events are implicitly connected and sorted by their time of occurrence and form a log. Then one could correlate the life-changing events across multiple persons and investigate surrounding events or inspect similar phases of life and compare the events they begin and end with. From a geological

#### 4. INCORR: INTERACTIVE DATA-DRIVEN CORRELATION PANELS FOR DIGITAL OUTCROP ANALYSIS

---

perspective, it is obvious that correlation panels can be used to represent time, because they already encode temporal aspects. The elevation of contacts and strata marks their occurrence in time, while the rank of a contact translates to the frequency of occurrence. For instance, the red contacts shown in Figure 4.1 represent a change that may happen once every 100 to 1000 years, while the changes shown as orange contacts can happen once in a week or a month.

# Conclusion

## 5.1 Summary

In this thesis we discussed and demonstrated tight integrations of geometric visualizations and attribute visualizations, building a bridge between visual analysis methods and 3D real-time rendering. In Chapter 1, we defined the dataset type of multivariate geometric data followed by specifying multivariate geometric tasks, which we based on existing task taxonomies concerned with spatial aspects of data. We then illustrated challenges stemming from the geometric nature of the data and tasks at hand and proposed five design goals to support visualization designers in alleviating these challenges. We further used these design goals to structure related work and illustrated that successful visualization solutions dealing with multivariate geometric data originate from domain-specific solutions. The number of different domains and domain-specific solutions show the relevance of multivariate geometric data and the necessity of adequate visualization methods, while the theoretical frameworks and models do not sufficiently address these data or tasks. With this thesis we address this gap in the literature by following a bottom up methodology, conducting three design studies addressing domain-specific problems from which we abstract general insights to refine visualization guidelines [SMM12].

Chapter 2 begins with a design study supporting experts in tunnel crack analysis in the context of a tunnel maintenance scenario. Similar to the tunnel use case in Section 1.1, we used the tunnel crack data and the high-resolution tunnel surface geometry as an illustrative example to highlight integration challenges inherent to 3D visualizations. We provided a variety of domain-specific implementations, as for instance a guided navigation method for locating tunnel cracks or a virtual X-ray technique to deal with occlusion, allowing domain experts to effectively analyze geometric and multivariate aspects simultaneously. As a next step, we generalized from these concrete implementations and derived VISAR, a methodological framework comprising the components ‘Localization’, ‘Overview & Detail’, ‘Visual Discrimination’, ‘Occlusion Handling’, ‘Visual Abstraction’, and ‘Similarity-based Analysis’. We also provided guidelines for each component

to support visualization designers in creating effective solutions in the presence of multivariate geometric data and tasks.

In Chapter 3, we presented Vis-A-Ware, a design study to support analysts in quantifying and analyzing the visual impact of planned buildings in an urban environment. We developed linked views for a quantitative and a qualitative analysis of multiple output values of variations with respect to measurement probes. We illustrated the applicability of our approach based on a use case scenario involving complex visibility questions. Qualitative feedback from experts of three different fields that are involved in the urban planning process confirmed that Vis-A-Ware enables a more holistic and efficient decision-making process, enriching a purely qualitative process with quantitative features.

We dedicated Chapter 4 to InCorr, a design study of translating the static geological illustration of correlation panels for outcrop analysis to an interactive solution. It comprises a 3D logging tool and an interactive data-driven correlation panel, the InCorrPanel. We integrated both with an existing outcrop interpretation software and provided a vector graphics export to integrate the results with the established workflow of geologists. We demonstrated InCorr's functionality through a use case and validated our results against a manually created illustration based on identical outcrop interpretations. Furthermore, we collected qualitative feedback from three geologists. The evaluation indicates that InCorr is a powerful and useful tool and can significantly improve geological analysis workflows through data propagation, and reduce time and effort of illustrating logs and correlation panels manually.

## 5.2 Outlook

We have demonstrated that multivariate geometric data and respective tasks are relevant in numerous domains and that there is a continuous need for addressing the intricacies of this problem space. Based on the work discussed in this thesis we identify the following directions for future research.

**Extend VISAR.** We successfully employed the VISAR methodology for integrating geometric and attribute views in a series of projects and published works, including but not limited to the works presented in this thesis. When looking at the framework illustrated in Figure 2.8, most components only apply to the geometric visualization. This is due to the fact, that visual perception tasks or `locate`, `identify`, and `compare` are most likely to fail in the geometric view. However, we realized, that the visual feedback of, for instance, scatter plots reacting to a selection in the geometric view is often not sufficient to direct a user's visual attention. This issue could be alleviated using visual links [WPL<sup>+</sup>10] connecting geometric representations and attribute representations. Another limitation of the VISAR methodology is that it only addresses existing data therefore only supports the `consuming` tasks of our taxonomy. It would be interesting to revisit the methodology to derive components supporting the `producing` tasks of `annotate` and `derive`. For instance, the quantification of visual impact in Chapter 3 is based on a simple pixel count, while it took several design iterations to derive intuitive values for visibility occlusion and sky occlusion. User feedback indicated, that we did not succeed in providing an intuitive quantification for openness occlusion and that the other metrics could

be improved by including the shape of candidate buildings into the impact value. In some domains quantification for certain geometric and physical phenomena are already provided, as for instance uniformity and glare rating in indoor lighting design [SOL<sup>+</sup>15, WSL<sup>+</sup>19]. However, visualization designers often need to research new quantifications or improve on existing ones and would benefit from a methodological approach and guidelines.

**Common Coordinate Axes.** In all of the presented design studies we chose the approach of faceting different aspects of our data in multiple views, that is, geometric representations in geometric views and multivariate aspects in attribute views. But, some of the data aspects we visualize in attribute views do have a spatial meaning, such as Sv stationing (i.e., metric along the tunnel axis) of a tunnel crack or the thickness of a rock layer. Recalling the aggregation view in Figure 2.3d, we used the Sv coordinate as the horizontal axis implicitly arranging moisture distributions of cracks along the tunnel. While the horizontal axis in the InCorrPanel is misleading for people not familiar with this kind of representation, the vertical axis of each log does represent the course of the 3D log in the geometric view in an undistorted form. These shared coordinate axes are very domain specific, and stationing metrics are typical for linear structures, such as tunnels, roads, hiking trails or rivers. However, it would be interesting to explore examples from other domains, as for instance the pathway of a biopsy needle [HMP<sup>+</sup>12] or ‘stationing’ along straightened blood vessels [AH11], and derive a more general visualization paradigm for generating such one dimensional shared coordinate.

**State-of-the-Art Report.** Based on our definition of data, tasks, and challenges we derived design goals. We could use the sophistication concerning the fulfillment of each design goal as an axis in a design space. Since we already employed this system to structure our related work section it would be interesting to collect a comprehensive overview of multivariate geometric data and respective integrated solutions in the form of a state-of-the-art report.

**Multivariate time-dependent Geometric Data.** This thesis and the discussed design studies are concerned with multivariate geometric data. When inspecting a tunnel, evaluating candidates for urban development, or interpreting 3D rock outcrops domain expert interactions are limited to static scenarios. For instance, assuming that a tunnel surface and the cracks in it are only digitally captured biannually it is usually sufficient to just inspect the two latest temporal instances of a tunnel. However, the advance of new acquisition techniques and platforms steadily shorten the acquisition intervals, slowly converging to a steady stream of updated geometric representations. While this trend currently poses a challenge to processing pipelines, data bases, and 3D rendering techniques, we see the necessity to extend the concept of multivariate geometric data to multivariate time-dependent geometric data. This entails an extended set of tasks, design goals, and a new methodological framework providing guidelines to address newly emerging challenges.

With this overview of interesting future prospects of multivariate geometric data and the integration of respective visualizations, we conclude this work. We hope that this thesis, by extending the existing body of knowledge on the integration of geometric and attribute visualizations, shapes the course of future design studies, visualization solutions, and visualization research.



# List of Figures

1.1	(a) 3D geometric view showing tunnel cracks as 3D polylines on high-resolution textured tunnel surface reconstruction. (b) Detail view on selected tunnel crack in red. (c) Attribute data of tunnel cracks in the form of rows in a table, with the selected crack highlighted in red. Adapted from [OSP <sup>+</sup> 16] . . . . .	2
1.2	(a) Textured 3D reconstruction derived from tunnel surface acquisition by laser scanners and high-resolution cameras. (c) Camera frustum of the acquisition process that contains 3D query point (red) allows users to retrieve (b) the respective unprocessed source image. <i>The data has been provided by Dibit Messtechnik GmbH.</i> . . . . .	3
1.3	Series of so-called tunnel faces as captured during the advance of the tunnel head. Geological annotations indicate contacts between rock layers. <i>The data has been provided by Dibit Messtechnik GmbH and has been annotated by Robert Barnes.</i>	4
1.4	Grey - Dataset types composed of core datatypes as presented in Munzner [Mun14]. Yellow - ‘Multivariate Geometry’ as new dataset type, differing from ‘Geometry’ by ‘Shape’ and ‘Attributes’. . . . .	5
1.5	Mapping of typical patterns in the visual analysis of scientific data [ODH <sup>+</sup> 07] and visual perception tasks [ET08] onto Munzner’s [Mun14] ‘Why: Task Abstraction’ to define tasks for consuming multivariate geometric data (T1, T2, T3, and T4). New multivariate geometric data can be produced either through annotating or deriving (T5 and T6). . . . .	6
1.6	Brushing attribute values in scatter plot (b). Revealing ‘infarction core’ in brain in (a) which is color coded by attribute value. Smooth brushing in (d) reveals wider tissue with infarction risk. From [ODH <sup>+</sup> 07], simplified caption. . . . .	9
1.7	Brain is rotated in (e) for a better overview on infarction in 3D. Smooth brushing in (f) used to convey the shape of the brain. Smooth brushing of other attributes (h) yields similar result in (g) compared to Figure 1.6c. From [ODH <sup>+</sup> 07], simplified caption. . . . .	10
1.8	(a) 2D scatter plot showing orientation vs. length: linked selection of cracks, which are oriented along the tunnel direction, i.e., 0°, and which are longer than 8 m. (b) The scene rendered from the current camera position showing only some of the selected cracks (red) (c) Some cracks are partially outside, others are completely outside the view frustum. (d) Cracks occluded by the first tunnel wall are not visible from the current camera position. Adapted from Ortner et al. [OSP <sup>+</sup> 16]. . . . .	11

1.9	Comparing lighting parametrizations according to different metrics in the LiteVis workspace. (a) The geometric view displays a false color rendering encoding illumination quality. For each table, (b) floating annotations show a binned histogram of illumination values from selected simulation runs. (c) The ranking view shows a ranking of simulated lighting parametrizations according to user defined importance values for measurement surfaces and result indicators. From [SOL <sup>+</sup> 15], adapted caption. . . . .	15
2.1	(a) Reconstructed mesh of the tunnel surface. Tunnels provide an additional metric coordinate system consisting of the value Sv running along the tunnels axis and B running perpendicular to it. Each crack comprises (b) its geometric representation, a 3D polyline, and (c) its multivariate attributes, such as length, width, moisture, and orientation. Red highlighting indicates the same crack across (a),(b), and (c). . . . .	22
2.2	(a) 2D scatter plot showing orientation vs. length: linked selection of cracks, which are oriented along the tunnel direction, i.e., 0°, and which are longer than 8 m. (b) The scene rendered from the current camera position showing only some of the selected cracks (red) (c) Some are partially outside, others are completely outside of the view frustum. (d) Cracks occluded by the first tunnel wall are not visible from the current camera position. . . . .	24
2.3	System overview: (a) the geometric view shows a textured mesh of the north and south tunnel tube including tunnel cracks as polylines. (b) The scatter plot shows cracks as dots with respect to orientation on the x-axis vs. length on the y-axis. (c) The parallel coordinates plot enables the comparison and identification of trends for many dimensions. (d) The aggregation plot shows the distributions of moisture values in the north and south tunnel tube. Histograms are grouped to 120 m long intervals of Sv. (e) Selections on the length and orientation axis are combined by a logical AND. (f) Peek-brushing over the histogram causes blue highlights on the corresponding cracks in (a), (b), and (c). . . . .	26
2.4	Visual Discrimination: Different visual variables are manipulated to emphasize the differences between peek-brushed, brushed, and context cracks. (a) Three different widths are used, while cracks in focus have a superimposed glow. (b) Using different widths and colors: blue, red, and yellow-green for peek-brushed, brushed, and other cracks, respectively. (c) Color is used for attribute mapping, while visual discrimination is only conveyed by width and glow. . . . .	28
2.5	Localization Transition: (a) Users brush a single crack in the scatter plot. (b) A localization transition is triggered (1) the camera turns to the selected crack, (2) the camera moves to the closest point on a computed orbit sphere (3) the camera rotates around the selected crack and ends up at the characteristic viewpoint inside the tunnel. (c) The scene rendered from the characteristic viewpoint. . . . .	29



2.6	(a) Users investigate the distribution of cracks, which are orthogonal to the tunnel direction. (b) The selection triggers an automatic camera transition to a user-defined overview viewpoint. (c) Clicking on individual cracks seen from the overview viewpoint issues a localization transition, an additional click transitions back to the overview. . . . .	30
2.7	(a) Brushed cracks are highlighted in red including a glow effect. The glow effect persists even if the respective cracks are occluded. (b) Detailed tunnel crack representations are replaced by point sprites to avoid clutter. . . . .	32
2.8	The VISAR framework is divided into two layers: the mirroring layer and the integration layer. The mirroring layer contains simple coordinations, such as ‘Selection’, ‘Peek-Selection’, and ‘Color’. The components of the integration layer are concerned with more complex coordinations that facilitate the visual perception tasks (see section 2.3): ‘Guided Navigation’, ‘Enhanced Geometric Rendering’, and ‘Similarity-based Analysis’. . . . .	35
3.1	Data Model. (a) Evaluations are represented by an n*m matrix of candidates and viewpoints. Each combination holds a series of visual impact metric (VIM) values. (b) This representation is transformed into a lookup table either governed by viewpoints (viewpoint-major) or by candidates (candidate-major). . . . .	44
3.2	False color image evaluation of (a) existing situation, (b) for each candidate. (c) Four metrics to quantify visual impact: landmark occlusion, visibility, openness, sky occlusion. (d) ‘All candidates’ evaluation for normalization of metrics (except for the landmark occlusion metric). (e) Ratio of ‘all candidates’ to image size: large, medium, small. . . . .	45
3.3	20 exemplary candidate buildings of different sizes and shapes identified by letters from A to T. . . . .	48
3.4	(a) Bird’s-eye and (b) close up view, both showing the project area near the railway station with all building candidates and a landmark (orange). Further, we show exemplary viewpoints at street level and on the train tracks, as well as a line of sight. . . . .	49
3.5	An overview of Vis-A-Ware, a tool supporting 3D visibility analysis in urban planning. It features the following coordinated views: (a) a 3D spatial view to investigate spatial relations between planned structures and viewpoints, (b) the transposable ranking view to analyze derived visibility metrics, and the filmstrip (c) for a qualitative comparison of single-building renderings. (d) 3D representations of the candidate buildings are blended together, hovering in linked views renders the respective candidate as solid. . . . .	50
		99

3.6	Transposable Ranking View (TRV). (a) Detail bar charts showing individual VIM values of each candidate. (b) Stacked bar charts as a compact version of VIM distributions. (c) Linked peek-brushing shows detail-on-demand tooltip and a candidate's position within all other VIM distributions. (d) Rows are ranked by distribution scores of the selected VIM - in this case candidate visibility. (e) The arrow-icon loads the perspective from a viewpoint into the 3D view, while the film icon loads respective screenshot data into the filmstrip. (f) Each category of viewpoints is collapsible and features a heatmap-like aggregation of VIM scores. (g) Viewpoints are grouped by their coverage value, which is indicated by the size of the black circles. . . . .	52
3.7	Focus, Filter, and Transpose workflow for reducing candidate and viewpoint sets. (a) Twenty candidates impact six viewpoints. Selection of high class visibility candidates in viewpoint <i>vp50</i> . (b) Higher bars encode selected candidates (in focus). Selection of additional high class visibility candidates in viewpoint <i>vp48</i> . (c) The focus set covers all candidates exhibiting high visibility values. (d) Candidates are filtered to the focus set, dropping all candidates, which do not exhibit high visibility values in any viewpoint. (e) Transposed view: candidates are now rows, the selection now affects viewpoints highly impacted by candidate <i>T</i> . (f) Selection is extended by other highly impacted viewpoints. (g) All highly impacted viewpoints are now part of the focus set. (h) Set is filtered to the focus set. (i) Transpose to see the set reduced to six candidates in four viewpoints. . . . .	54
3.8	Q1. Which candidates do cover a landmark and how strong is the occlusion? (a) Users rank by landmark occlusion and (b) focus on very high and high impact values. (c) All very high and high classes are part of the focus set and (d) filtering reduces the set of candidates to five. (e) To investigate top scoring candidates <i>F</i> and <i>T</i> in detail, the TRV is transposed to candidate-major mode. (f) Users focus on viewpoints exhibiting very high impact values, which are then assessed through (g) a qualitative comparison in the filmstrip. . . . .	57
3.9	Q2. Which candidates are most prominent as seen from the railway? (a) A spatial overview depicts the distribution of the visibility impact as seen from the railway. (b) Viewpoints are sorted by candidate visibility. (c) Candidates and viewpoints are reduced through multiple focus, filter, and transpose interactions as illustrated in Figure 3.7. (d) Peek-brushing of the highest scoring candidates reveals their scores within the distributions. (e) The remaining candidate-viewpoint combinations can be compared qualitatively in the filmstrip. . . . .	59

3.10	Q3: Which candidates have the highest impact on sky and openness in the vicinity of the project area? (a) Users rank viewpoints by the openness metric and focus on (i.e., select) very high values of openness occlusion. (b) The focus set now contains all very high openness values ignoring <i>vp63</i> . (c) Then users rank by sky occlusion and select high values. (d) The focus set contains all high sky occlusion values. (e) Now users filter candidates, which are not in focus. (f) Transposing the view lets the users investigate the remaining six candidates. (g) They again focus on viewpoints exhibiting very high openness occlusion and high sky occlusion and (h) filter to the focus set. (i) Sorting by either openness occlusion or sky occlusion yields the top ranked candidates <i>O</i> and <i>T</i> respectively. Hovering candidate rows and viewpoints shows detail bar charts and results in the highlighting of buildings and viewpoint circles in the spatial view. . . . .	61
4.1	(a) Map of the Hanksville-Burpee Dinosaur Quarry campaign with its four outcrops along the canyon. (b) 3D triangulated mesh as digital outcrop model (DOM) of Sol2 with interpretation showing bedsets and units and their respective contacts. (c) Cross bed measurements, which determine a layer's deposition direction indicated by dip-and-strike disks. . . . .	69
4.2	Annotated 3D outcrop model, followed by a sophisticated and a simplified log representation, both illustrated manually with a vector drawing tool. (a) Vertical axis encodes true thickness of strata. (b) Rose diagrams show the distribution of dipping orientations. (c) Individual styles and glyphs convey rock characteristics. (d) Dashed lines convey uncertainty. (e) Horizontal axis encodes grain sizes logarithmically, directly relating to rock types. . . . .	73
4.3	The apparent thickness of a stratum observed at an outcrop is often misleading. The rock may be broken off at a slanted angle or the stratum itself may be tilted. A dip-and-strike measurement is essential to determine a stratum's orientation and compute its true thickness. . . . .	74
4.4	Logs arranged in a manually illustrated correlation panel with (a) converging contacts, (b) spatial distances between logs along geographic direction, and (c) a fossil layer used as leveling horizon. . . . .	75
4.5	(a) Data model: tree of strata and their bounding contacts. (b) InCorrPanel log: (c) secondary log shows a tree cut at depth 1, while the primary log shows the leaf cut. (d) Rose diagrams aggregate cross bed orientations. (e) Dashed border conveys uncertain an rock type. . . . .	79
4.6	System overview of InCorr: (a) Outcrop View, showing a digital outcrop model interpreted in 3D by a geologist and annotated with the vertical 3D logging tool indicating layer thicknesses. (b) InCorrPanel, showing logs for two outcrops created directly from annotation data mimicking manual illustrations. (c) GUI to assign rock types to rock layers. . . . .	80
4.7	To create a 3D log, users (a) first pick a reference plane, (b) then they connect the contacts to form (c) a 3D log resulting in an initially (d) empty log. . . . .	83

4.8	Users assign rock types to the empty log via the (a) rock types list, resulting in (b) a log encoding rock categories and grain sizes. To complete the log, geologists (c) assign cross bed measurements, which are aggregated and encoded as (d) rose diagrams. . . . .	84
4.9	Outcrops Sol 2–5 of the HBDQ dataset, each with a 3D log created with InCorr and the resulting 2D log from the InCorrPanel. . . . .	87
4.10	(a) Correlation panel created with InCorr based on outcrop interpretation data with (b,c,d) manual illustrations with detail lenses (1-4) for comparison. (e) All logs are aligned to a common baseline correlation at the top. (f) No vertical exaggeration of strata, (g) no indication of buried regions, and no visible connection of non-adjacent logs. (h) Rulers showing spatial distance between logs. . . . .	88

# List of Tables

4.1	Summary of the most important geological terms. . . . .	77
-----	---	----



# Bibliography

- [AA03] Natalia Andrienko and Gennady Andrienko. Informed spatial decisions through coordinated views. *Information Visualization*, 2(4):270–285, 2003.
- [AAB<sup>+</sup>10] Gennady Andrienko, Natalia Andrienko, Sebastian Bremm, Tobias Schreck, Tatiana Von Landesberger, Peter Bak, and Daniel Keim. Space-in-time and time-in-space self-organizing maps for exploring spatiotemporal patterns. In *Computer Graphics Forum*, volume 29, pages 913–922. Wiley Online Library, 2010.
- [AH11] Paolo Angelelli and Helwig Hauser. Straightening tubular flow for side-by-side visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2063–2070, 2011.
- [AMHH11] Tomas Akenine-Möller, Eric Haines, and Naty Hoffman. *Real-time rendering*. CRC Press, 2011.
- [BBD05] Henrik Buchholz, Johannes Bohnet, and Jurgen Döllner. Smart and physically-based navigation in 3d geovirtual environments. In *Proceedings of the International Conference on Information Visualisation*, pages 629–635. IEEE, 2005.
- [BGT<sup>+</sup>18] Robert Barnes, Sanjeev Gupta, Christoph Traxler, Thomas Ortner, Arnold Bauer, Gerd Hesina, Gerhard Paar, Ben Huber, Kathrin Juhart, Laura Fritz, Bernhard Nauschnegg, Jan-Peter Muller, and Yu Tao. Geological analysis of Martian rover-derived digital outcrop models using the 3-d visualization tool, planetary robotics 3-d viewer PRo3d. *Earth and Space Science*, 5(7):285–307, 2018.
- [BM10] Stefan Bruckner and Torsten Möller. Result-driven exploration of simulation parameter spaces for visual effects design. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1468–1476, 2010.
- [BM13] Matthew Brehmer and Tamara Munzner. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–2385, December 2013.
- [BOH11] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. D<sup>3</sup> data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011.

- [BP10a] Wolfgang Berger and Harald Piringer. Interactive visual analysis of multiobjective optimizations. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology*, pages 215–216. IEEE, 2010.
- [BP10b] Wolfgang Berger and Harald Piringer. Peek brush: A high-speed lightweight ad-hoc selection for multiple coordinated views. In *Proceedings of the IEEE Conference on Information Visualization*, pages 140–145, 2010.
- [BPF11] Wolfgang Berger, Harald Piringer, Peter Filzmoser, and Eduard Gröller. Uncertainty-aware exploration of continuous parameter spaces using multivariate prediction. In *Computer Graphics Forum*, volume 30, pages 911–920, 2011.
- [BRN<sup>+</sup>19] Simon J. Buckley, Kari Ringdal, Nicole Naumann, Benjamin Dolva, Tobias H. Kurz, John A. Howell, and Thomas J.B. Dewez. Lime: Software for 3-d visualization, interpretation, and communication of virtual geoscience models. *Geosphere*, 15(1):222–235, 2019.
- [BRR03] Nadia Boukhelifa, Jonathan C. Roberts, and Peter J. Rodgers. A coordination model for exploratory multiview visualization. In *Proceedings of the International Conference on Coordinated and Multiple Views in Exploratory Visualization*, pages 76–85, 2003.
- [BSFL<sup>+</sup>12] Martin Brunnhuber, Helmut Schrom-Feiertag, Christian Luksch, Thomas Matyus, and Gerd Hesina. Bridging the gap between visual exploration and agent-based pedestrian simulation in a virtual environment. In *Proceedings of the 18th ACM symposium on Virtual reality software and technology*, pages 9–16. ACM, 2012.
- [BWPP04] Jiří Bittner, Michael Wimmer, Harald Piringer, and Werner Purgathofer. Coherent hierarchical culling: Hardware occlusion queries made useful. In *Computer Graphics Forum*, volume 23, pages 615–624. Wiley Online Library, 2004.
- [BWW<sup>+</sup>08] Thomas Butkiewicz, Dou Wenwen, Zachary Wartell, William Ribarsky, and Remco Chang. Multi-Focused Geospatial Analysis Using Probes. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1165–1172, 2008.
- [CBZ<sup>+</sup>06] Remco Chang, Thomas Butkiewicz, Caroline Ziemkiewicz, Zachary Wartell, Nancy Pollard, and William Ribarsky. Hierarchical simplification of city models to maintain urban legibility. In *Proceedings of ACM SIGGRAPH*, volume 6, page 130, 2006.
- [CBZ<sup>+</sup>08] Remco Chang, Thomas Butkiewicz, Caroline Ziemkiewicz, Zachary Wartell, Nancy Pollard, and William Ribarsky. Legible simplification of textured urban models. *IEEE Computer Graphics and Applications*, 28(3):27–36, 2008.
- [CKS<sup>+</sup>15] Daniel Cornel, Artem Konev, Bernhard Sadransky, Zsolt Horvath, Eduard Gröller, and Jürgen Waser. Visualization of object-centered vulnerability to possible flood hazards. In *Computer Graphics Forum*, volume 34, pages 331–340. Wiley Online Library, 2015.



- [Clo] CloudCompare. 3D point cloud and mesh processing software. <https://cloudcompare.org>. Accessed: 2020-04-28.
- [CWK<sup>+</sup>07] Remco Chang, Ginette Wessel, Robert Kosara, Eric Sauda, and William Ribarsky. Legible Cities: Focus-Dependent Multi-Resolution Visualization of Urban Relationships. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1169–1175, 2007.
- [DFL<sup>+</sup>15] Harish Doraiswamy, Nivan Ferreira, Marcos Lage, Huy Vo, Luc Wilson, Heidi Werner, Muchan Park, and Cláudio T. Silva. Topology-based catalogue exploration framework for identifying view-enhanced tower designs. *ACM Transactions on Graphics (TOG)*, 34(6):230, 2015.
- [ESA] ESA. Searching for signs of life on Mars. <http://exploration.esa.int/mars/43608-life-on-mars/>. Accessed: 2020-04-28.
- [Esr] Esri Geospatial Cloud. ArcGIS - the mapping and analytics platform. <https://www.esri.com/en-us/arcgis/about-arcgis/overview>. Accessed: 2020-04-28.
- [ET08] Niklas Elmqvist and Philippas Tsigas. A taxonomy of 3d occlusion management for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(5):1095–1109, 2008.
- [Fer04] Randima Fernando. *GPU Gems: Programming Techniques, Tips and Tricks for Real-Time Graphics*. Pearson Higher Education, 2004.
- [Fer05] Randima Fernando. Percentage-closer soft shadows. In *ACM SIGGRAPH 2005 Sketches*, page 35. 2005.
- [FH09] Raphael Fuchs and Helwig Hauser. Visualization of multi-variate scientific data. In *Computer Graphics Forum*, volume 28, pages 1670–1690. Wiley Online Library, 2009.
- [FLD<sup>+</sup>15] Nivan Ferreira, Marcos Lage, Harish Doraiswamy, Huy Vo, Luc Wilson, Heidi Werner, Muchan Park, and Cláudio T. Silva. Urbane: A 3d framework to support data driven decision making in urban development. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 97 – 104, 2015.
- [Gao16] Yang Gao. *Contemporary Planetary Robotics: An Approach Toward Autonomous Systems*. John Wiley & Sons, 2016.
- [Geo] Geoweb3d. Geoweb3d. <http://www.geoweb3d.com/>. Accessed: 2020-11-08.
- [GGM<sup>+</sup>15] John P. Grotzinger, Sanjeev Gupta, Michael C. Malin, David M. Rubin, Jürgen Schieber, Kirsten Siebach, Dawn Y. Sumner, Kathryn M. Stack, Ashwin R. Vasavada, Raymond E. Arvidson, et al. Deposition, exhumation, and paleoclimate of an ancient lake deposit, gale crater, mars. *Science*, 350(6257):aac7575, 2015.

- [GLG<sup>+</sup>13] Samuel Gratzl, Alexander Lex, Nils Gehlenborg, Hanspeter Pfister, and Marc Streit. Lineup: Visual analysis of multi-attribute rankings. *Visualization and Computer Graphics, IEEE Transactions on*, 19(12):2277–2286, 2013.
- [GRA20] GRASS Development Team. *Geographic Resources Analysis Support System (GRASS GIS) Software*. Open Source Geospatial Foundation, USA, 2020.
- [Gre] Marshall Greenblatt. Cef. <https://bitbucket.org/chromiumembedded/cef>. Accessed: 2015-06-30.
- [GRW<sup>+</sup>00] Donna L Gresh, Bernice E Rogowitz, Raimond L Winslow, David F Scollan, and Christina K Yung. WEAVE: a system for visually linking 3-D and statistical visualizations, applied to cardiac simulation and measurement data. In *Proceedings of the Conference on Visualization*, pages 489–492, 2000.
- [Hau06] Helwig Hauser. Generalizing focus+ context visualization. In *Scientific visualization: The visual extraction of knowledge from data*, pages 305–327. Springer, 2006.
- [HB03] Mark Harrower and Cynthia A Brewer. Colorbrewer. org: an online tool for selecting colour schemes for maps. *The Cartographic Journal*, 40(1):27–37, 2003.
- [HBG<sup>+</sup>11] Thomas Höllt, Johanna Beyer, Fritz Gschwantner, Philipp Muigg, Helmut Doleisch, Gabor Heinemann, and Markus Hadwiger. Interactive seismic interpretation with piecewise global energy minimization. In *2011 IEEE Pacific Visualization Symposium*, pages 59–66. IEEE, 2011.
- [HFG<sup>+</sup>12] Thomas Höllt, Wolfgang Freiler, Fritz Gschwantner, Helmut Doleisch, Gabor Heinemann, and Markus Hadwiger. SeiVis: An interactive visual subsurface modeling application. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2226–2235, 2012.
- [HGE<sup>+</sup>11] Alexander G. Hayes, John P. Grotzinger, Lauren A. Edgar, Steven W. Squyres, Wesley A. Watters, and Jasha Sohl-Dickstein. Reconstruction of eolian bed forms and paleocurrents from cross-bedded strata at Victoria crater, Meridiani Planum, Mars. *Journal of Geophysical Research: Planets*, 116(E7), 2011.
- [HGS<sup>+</sup>11] Gary J. Hampson, Royhan M. Gani, Kathryn E. Sharman, Nawazish Irfan, and Bryan Bracken. Along-strike and down-dip variations in shallow-marine sequence stratigraphic architecture: Upper cretaceous star point sandstone, Wasatch plateau, central Utah, U.S.A. *Journal of Sedimentary Research*, 81(3):159–184, March 2011.
- [HGWR07] David Hodgetts, Robert L. Gawthorpe, Paul Wilson, and Frank Rarity. Integrating digital and traditional field techniques using virtual reality geological studio (vrgs). In *69th EAGE Conference and Exhibition incorporating SPE EUROPEC 2007*, pages cp–27. European Association of Geoscientists & Engineers, 2007.

- [HMP<sup>+</sup>12] Paul-Corneliu Herghelegiu, Vasile Manta, Radu Perin, Stefan Bruckner, and Eduard Gröller. Biopsy planner–visual analysis for needle pathway planning in deep seated brain tumor biopsy. In *Computer Graphics Forum*, volume 31, pages 1085–1094. Wiley Online Library, 2012.
- [JDL09] Radu Jianu, Cagatay Demiralp, and David H. Laidlaw. Exploring 3d DTI fiber tracts with linked 2d representations. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1449–1456, 2009.
- [JE12] Waqas Javed and Niklas Elmquist. Exploring the design space of composite visualization. In *Pacific Visualization Symposium (PacificVis), 2012 IEEE*, pages 1–8, 2012.
- [JJF07] Mikael Jern, Sara Johansson, Jimmy Johansson, and Johan Franzen. The gav toolkit for multiple linked views. In *Proceedings of the International Conference on Coordinated and Multiple Views in Exploratory Visualization*, pages 85–97, 2007.
- [JKKS20] Stefan Jänicke, Pawandeep Kaur, Pawel Kuzmicki, and Johanna Schmidt. Participatory visualization design as an approach to minimize the gap between research and application. In *Proceedings of the Workshop on the Gap between Visualization Research and Visualization Software (VisGap)*, 2020.
- [JPJW16] Richard R. Jones, Mark A. Pearce, Carl Jacquemyn, and Francesca E. Watson. Robust best-fit planes from geospatial data. *Geosphere*, 12(1):196–202, 2016.
- [JZ10] Sudabe Jafari and Narges Zaredar. Land suitability analysis using multi attribute decision making approach. *International journal of environmental science and development*, 1(5):441–445, 2010.
- [Kei02] Daniel A Keim. Information visualization and visual data mining. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):1–8, 2002.
- [KH13] Johannes Kehler and Helwig Hauser. Visualization and Visual Analysis of Multifaceted Scientific Data: A Survey. *IEEE Transactions on Visualization and Computer Graphics*, 19(3):495–513, 2013.
- [KS64] William C. Krumbein and Laurence L. Sloss. Stratigraphy and sedimentation. *Acta Crystallographica*, 17(8):1090–1090, August 1964.
- [LNP<sup>+</sup>13] Endre M Lidal, Mattia Natali, Daniel Patel, Helwig Hauser, and Ivan Viola. Geological storytelling. *Computers & Graphics*, 37(5):445–459, 2013.
- [LRC<sup>+</sup>03] David Luebke, Martin Reddy, Jonathan D. Cohen, Amitabh Varshney, Benjamin Watson, and Robert Huebner. *Level of detail for 3D graphics*. Morgan Kaufmann, 2003.

- [LTH<sup>+</sup>13] Christian Luksch, Robert F. Tobler, Ralf Habel, Michael Schwärzler, and Michael Wimmer. Fast light-map computation with virtual polygon lights. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, I3D'13*, pages 87–94, New York, NY, USA, 2013. ACM.
- [MDL<sup>+</sup>18] Fabio Miranda, Harish Doraiswamy, Marcos Lage, Luc Wilson, Mondrian Hsieh, and Cláudio T. Silva. Shadow accrual maps: Efficient accumulation of city-scale shadows over time. *IEEE Transactions on Visualization and Computer Graphics*, 25(3):1559–1574, 2018.
- [MFGH08] Kresimir Matkovic, Wolfgang Freiler, Denis Gracanin, and Helwig Hauser. Comvis: A coordinated multiple views system for prototyping new visualization technology. In *Proceedings of the IEEE Conference on Information Visualization*, pages 215–220, 2008.
- [MR] Microsoft-Research. F#. <http://fsharp.org/>. Accessed: 2015-06-30.
- [Mun14] Tamara Munzner. *Visualization analysis and design*. CRC press, 2014.
- [MvGW11] Malgorzata A. Migut, Jan C. van Gemert, and Marcel Worring. Interactive decision making using dissimilarity to visually represented prototypes. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 141–149, 2011.
- [NAS] NASA. Mars 2020 mission overview. <https://mars.nasa.gov/mars2020/mission/overview/>. Accessed: 2020-04-28.
- [NDH<sup>+</sup>18] Paul Ryan Nesbit, Paul R. Durkin, Christopher H. Hugenholtz, Stephen M. Hubbard, and Maja Kucharczyk. 3-d stratigraphic mapping using a digital outcrop model derived from uav images and structure-from-motion photogrammetry. *Geosphere*, 14(6):2469–2486, 2018.
- [NLP<sup>+</sup>13] Mattia Natali, Endre M Lidal, Julius Parulek, Ivan Viola, and Daniel Patel. Modeling terrains and subsurface geology. In *Eurographics (STARs)*, pages 155–173, 2013.
- [NM14] Jae-Ho Nah and Dinesh Manocha. Sato: Surface area traversal order for shadow ray tracing. In *Computer Graphics Forum*, volume 33, pages 167–177. Wiley Online Library, 2014.
- [NS00] Chris North and Ben Shneiderman. Snap-together visualization: a user interface for coordinating visualizations via relational schemata. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, pages 128–135, 2000.
- [ODH<sup>+</sup>07] Steffen Oeltze, Helmut Doleisch, Helwig Hauser, Philipp Muigg, and Bernhard Preim. Interactive visual analysis of perfusion data. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1392–1399, 2007.

- [OPH<sup>+</sup>10] Thomas Ortner, Gerhard Paar, Gerd Hesina, Robert F Tobler, and Bernhard Nauschnegg. Towards true underground infrastructure surface documentation. *Proceedings of Real CORP*, 2010.
- [OSP<sup>+</sup>16] Thomas Ortner, Johannes Sorger, Harald Piringer, Gerd Hesina, and Eduard Gröller. Visual analytics and rendering for tunnel crack analysis. *The Visual Computer*, 32(6-8):859–869, 2016.
- [OSS<sup>+</sup>16] Thomas Ortner, Johannes Sorger, Harald Steinlechner, Gerd Hesina, Harald Piringer, and Eduard Gröller. Vis-a-ware: Integrating spatial and non-spatial visualization for visibility-aware urban planning. *IEEE Transactions on Visualization and Computer Graphics*, 23(2):1139–1151, 2016.
- [OTP<sup>+</sup>19] Thomas Ortner, Christoph Traxler, Harald Piringer, Laura Fritz, Maria Schimkowitz, Gerhard Paar, Gerhard Triebnig, Fabian Schindler, and Bernhard Nauschnegg. Minerva: A 3d GIS and visual analysis framework. In *Proceedings of 15th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA 2019)*, 2019.
- [OWN<sup>+</sup>21] Thomas Ortner, Andreas Walch, Rebecca Nowak, Robert Barnes, Thomas Höllt, and Eduard Gröller. InCorr: Interactive data-driven correlation panels for digital outcrop analysis. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):755–764, 2021.
- [PBK10] Harald Piringer, Wolfgang Berger, and Jürgen Krasser. Hypermoval: Interactive visual validation of regression models for real-time simulation. In *Computer Graphics Forum*, volume 29, pages 983–992. Wiley Online Library, 2010.
- [PCPKS06] Gerhard Paar, Maria del Pilar Caballo-Perucha, Heiner Kontrus, and Oliver Sidla. Optical crack following on tunnel surfaces. In *SPIE Proceedings of Two- and Three-Dimensional Methods for Inspection and Metrology IV*, volume 6382, 2006.
- [PGT<sup>+</sup>08] Daniel Patel, Christopher Giertsen, John Thurmond, John Gjelberg, and Eduard Gröller. The seismic analyzer: Interpreting and illustrating 2d seismic data. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1571–1578, 2008.
- [PS96] Donald R. Prothero and Fred Schwab. *An introduction to sedimentary rocks and stratigraphy*. WH Freeman and Company, 1996.
- [QGI] QGIS Development Team. QGIS. <https://qgis.org/en/site/>. Accessed: 2020-10-25.
- [Rob07] Jonathan C. Roberts. State of the Art: Coordinated Multiple Views in Exploratory Visualization. In *Proceedings of the International Conference on Coordinated and Multiple Views in Exploratory Visualization*, pages 61–71, 2007.

- [RWF<sup>+</sup>13] Hrvoje Ribičić, Jürgen Waser, Raphael Fuchs, Günter Blöschl, and Eduard Gröller. Visual Analysis and Steering of Flooding Simulations. *IEEE Transactions on Visualization and Computer Graphics*, 19(6):1062–1075, 2013.
- [SBS<sup>+</sup>13] Johannes Sorger, Katja Bühler, Florian Schulze, Tianxiao Liu, and Barry Dickson. neuroMap - Interactive Graph Visualization of the Fruit Fly’s Neural Circuit. In *IEEE Symposium on Biological Data Visualization*, pages 73–80, 2013.
- [Sch] Schlumberger Information Solutions. Petrel seismic to simulation software. <https://www.software.slb.com/products/petrel/petrel-geophysics/seismic-interpretation>. Accessed: 2020-04-28.
- [Sch11] Hans-Jörg Schulz. Treevis. net: A tree visualization reference. *IEEE Computer Graphics and Applications*, 31(6):11–15, 2011.
- [SG15] Hiranya Sahoo and Nahid D. Gani. Creating three-dimensional channel bodies in LiDAR-integrated outcrop characterization: A new approach for improved stratigraphic analysis. *Geosphere*, 11(3):777–785, 2015.
- [SHB<sup>+</sup>14] Michael Sedlmair, Christoph Heinzl, Stefan Bruckner, Harald Piringer, and Torsten Möller. Visual parameter space analysis: A conceptual framework. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2161–2170, 2014.
- [SMM12] Michael Sedlmair, Miriah Meyer, and Tamara Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440, December 2012.
- [Sof13] Tableau Software. Tableau, 2013. Accessed: 2013-10-07.
- [SOL<sup>+</sup>15] Johannes Sorger, Thomas Ortner, Christian Luksch, Michael Schwärzler, Eduard Gröller, and Harald Piringer. Litevis: Integrated visualization for simulation-based decision support in lighting design. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):290–299, 2015.
- [SOP<sup>+</sup>15] Johannes Sorger, Thomas Ortner, Harald Piringer, Gerd Hesina, and Eduard Gröller. A taxonomy of integration techniques for spatial and non-spatial visualizations. In *Proceedings of the Eurographics Symposium on Vision, Modeling and Visualization*, pages 57–64, 2015.
- [SP19] David J. Sanderson and David C.P. Peacock. Making rose diagrams fit-for-purpose. *Earth-Science Reviews*, page 103055, 2019.
- [SPA<sup>+</sup>14] Johanna Schmidt, Reinhold Preiner, Thomas Auzinger, Michael Wimmer, Eduard Gröller, and Stefan Bruckner. YMCA - your mesh comparison application. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pages 153–162, 2014.

- [STKD12] Amir Semmo, Matthias Trapp, Jan Eric Kyprianidis, and Jürgen Döllner. Interactive visualization of generalized virtual 3d city models using level-of-abstraction transitions. In *Computer Graphics Forum*, volume 31, pages 885–894, 2012.
- [TBPD11] Matthias Trapp, Christian Beesk, Sebastian Pasewaldt, and Jürgen Döllner. Interactive rendering techniques for highlighting in 3d geovirtual environments. In *Advances in 3D Geo-Information Sciences*, pages 197–210. Springer, 2011.
- [The20] TheCaleydoTeam. Caleydo, 2020. Accessed: 2020-11-03.
- [TIB20] TIBCO. Spotfire, 2020. Accessed: 2020-11-03.
- [TM04] Melanie Tory and Torsten Möller. Rethinking Visualization: A High-Level Taxonomy. In *Proceedings of the IEEE Conference on Information Visualization*, pages 151–158, 2004.
- [TMH09] Suvi Tarkkanen, Kaisa Miettinen, and Jussi Hakanen. Interactive poster: Interactive multiobjective optimization—a new application area for visual analytics. In *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on*, pages 237–238. IEEE, 2009.
- [TN08] Christoph Traxler and Wolfgang Neubauer. The harris matrix composer - a new tool to manage archaeological stratigraphy. In *Archäologie und Computer-Kulturelles Erbe und Neue Technologien-Workshop*, volume 13, pages 3–5, 2008.
- [VFSG06] Ivan Viola, Miquel Feixas, Mateu Sbert, and Eduard Gröller. Importance-Driven Focus of Attention. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):933–940, 2006.
- [Vis] Visplore GmbH. Visplore. <https://visplore.com/>. Accessed: 2020-11-07.
- [vLHRFP09] Xavier M.T. van Lanen, David Hodgetts, Jonathan Redfern, and Ivan Fabuel-Perez. Applications of digital outcrop models: two fluvial case studies from the Triassic Wolfville Fm., Canada and Oukaimeden Sandstone Fm., Morocco. *Geological Journal*, 44(6):742–760, 2009.
- [VM02] Gokul Varadhan and Dinesh Manocha. Out-of-core rendering of massive geometric environments. In *IEEE Visualization, 2002. VIS 2002*, pages 69–76. IEEE, 2002.
- [VRVa] VRVis. Aardvark. <https://aardvarkians.com/>. Accessed: 2020-04-28.
- [VRVb] VRVis. PRo3D. <http://pro3d.space>. Accessed: 2020-04-28.
- [WBWK00] Michelle Q. Wang-Baldonado, Allison Woodruff, and Allan Kuchinsky. Guidelines for using multiple views in information visualization. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, pages 110–119, 2000.

- [Wea04] Chris Weaver. Building highly-coordinated visualizations in improvise. In *Proceedings of the IEEE Conference on Information Visualization*, pages 159–166, 2004.
- [WFR<sup>+</sup>10] Jürgen Waser, Raphael Fuchs, Hrvoje Ribicic, Benjamin Schindler, Günther Blöschl, and Eduard Gröller. World Lines. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1458–1467, 2010.
- [Wla18] Scott Wlaschin. *Domain Modeling Made Functional: Tackle Software Complexity with Domain-Driven Design and F#*. Pragmatic Bookshelf, 2018.
- [WPL<sup>+</sup>10] Manuela Waldner, Werner Puff, Alexander Lex, Marc Streit, and Dieter Schmalstieg. Visual links across applications. In *Proceedings of Graphics Interface*, pages 129–136, 2010.
- [WSL<sup>+</sup>19] Andreas Walch, Michael Schwärzler, Christian Luksch, Elmar Eisemann, and Theresia Gschwandtner. Lightguider: Guiding interactive lighting design using suggestions, provenance, and quality visualization. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):569–578, 2019.
- [YPL07] Perry P.-J. Yang, Simon Y. Putra, and Wenjing Li. Viewsphere: a GIS-based 3d visibility analysis for urban design evaluation. *Environment and Planning B: Planning and Design*, 34(6):971–992, 2007.
- [ZFA<sup>+</sup>14] Wei Zeng, Chi-Wing Fu, Stefan Müller Arisona, Alexander Erath, and Huamin Qu. Visualizing mobility of public transportation system. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1833–1842, 2014.
- [ZYM<sup>+</sup>14] Jiawan Zhang, E Yanli, Jing Ma, Yahui Zhao, Bingham Xu, Liting Sun, Jinyan Chen, and Xiaoru Yuan. Visual analysis of public utility service problems in a metropolis. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1843–1852, 2014.



# Thomas Ortner – Curriculum Vitae

---

<b>Address</b>	Ottakringerstraße 128/18, 1160, Vienna	<b>Mobile Phone</b>	+43 (676) 373 1883
<b>Date of Birth</b>	12 <sup>th</sup> of December 1983	<b>Email</b>	ortner.t@googlemail.com
<b>Nationality</b>	Austrian	<b>Homepage</b>	thomasortner.github.io

## Education

**2012-2021** Ph.D., TU Wien, Vienna, Austria.  
Doctoral Programme in Technical Sciences (*ongoing*)  
*Thesis: Tight Integration of Visual Analysis and 3D Real-Time Rendering*

**2007-2008** M.Sc., University of Hull, Kingston upon Hull, United Kingdom.  
Computer Graphics (*with distinction*)  
*Thesis: GPGPU Ray Tracing in Real Time*

**2006-2007** M.Sc., School of Applied Sciences Technikum Vienna, Vienna, Austria.  
Software Engineering (*with distinction*)  
*Thesis: Out-of-Core Rendering of Massive Point Clouds*

**2003-2006** B.Sc., School of Applied Sciences Technikum Vienna, Vienna.  
Computer Science

## Employment History

**2020 - present** VRVis Zentrum für Virtual Reality und Visualisierung Forschungs-GmbH. Vienna. Austria.  
*Head of Group*  
*Geospatial Visualization, Semantic Modelling, and Acquisition (GeoSMAQ)*. Planning and coordination of applied and strategic research efforts for multiple projects.

**2016 - present** VRVis Zentrum für Virtual Reality und Visualisierung Forschungs-GmbH. Vienna. Austria.  
*Project Manager & Researcher*  
*TSID – Tunnel Surface Inspection and Documentation*. Project planning, design and realization of features, and communication of results.  
*PRo3D – Planetary Robotics 3D Viewer*. Project planning, design and realization of features, and communication of results in a multitude of space related projects supporting PRo3D.

**2012-2016** VRVis Zentrum für Virtual Reality und Visualisierung Forschungs-GmbH. Vienna. Austria.  
*Software Engineer & Researcher*  
*PRo3D*. Research and development of PRo3D collaborating with planetary scientists.

**2012-2016** VRVis Zentrum für Virtual Reality und Visualisierung Forschungs-GmbH. Vienna. Austria.  
*Ph.D. Researcher*  
*VISAR – Visual Analytics and Rendering*. Fundamental research project concerned with the tight integration of 3D geometric and multivariate attribute data, by combining 3D real-time visualizations and information visualization views.

**2008-2012** VRVis Zentrum für Virtual Reality und Visualisierung Forschungs-GmbH. Vienna. Austria.  
*Software Engineer & Researcher*  
*Dibit7*. Development of the core visualization component within a tunnel documentation application (precursor of TSID).

## Skills

<i>Professional</i>	3D Real-Time Rendering, Information Visualization & Visual Analytics, Domain Driven Design Functional Programming
<i>Code</i>	F#, JavaScript/TypeScript, C#, C++
<i>Frameworks &amp; Platforms</i>	D3.js, DirectX/OpenGL, .net, WPF, CUDA
<i>Languages</i>	German ( <i>native</i> ), English ( <i>fluent</i> )
<i>Other</i>	Gender & Diversity Training, JIRA, Confluence, Scrum, GIT/SVN, LaTeX

## Publications (selected)

- (1) **Thomas Ortner**, Andreas Walch, Rebecca Nowak, Robert Barnes, Thomas Höllt, and Eduard Gröller. Incorr: Interactive data-driven correlation panels for digital outcrop analysis. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):755–764, 2021.
- (2) **Thomas Ortner**, Christoph Traxler, Harald Piringer, Laura Fritz, Maria Schimkowitz, Gerhard Paar, Gerhard Triebnig, Fabian Schindler, and Bernhard Nauschnegg. Minerva: A 3d GIS and visual analysis framework. In *Proceedings of 15th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA 2019)*, 2019.
- (3) Robert Barnes, Sanjeev Gupta, Christoph Traxler, **Thomas Ortner**, Arnold Bauer, Gerd Hesina, Gerhard Paar, Ben Huber, Kathrin Juhart, Laura Fritz, Bernhard Nauschnegg, Jan-Peter Muller, and Yu Tao. Geological analysis of Martian rover derived digital outcrop models using the 3D visualization tool, planetary robotics 3D viewer PRo3d. *Earth and Space Science*, 5(7):285–307, 2018.
- (4) **Thomas Ortner**, Johannes Sorger, Harald Steinlechner, Gerd Hesina, Harald Piringer, and Eduard Gröller. Vis-a-Ware: Integrating spatial and non-spatial visualization for visibility-aware urban planning. *IEEE Transactions on Visualization and Computer Graphics*, 23(2):1139–1151, 2016.
- (5) **Thomas Ortner**, Johannes Sorger, Harald Piringer, Gerd Hesina, and Eduard Gröller. Visual analytics and rendering for tunnel crack analysis. *The Visual Computer*, 32(6-8):859–869, 2016.
- (6) Johannes Sorger, **Thomas Ortner**, Christian Luksch, Michael Schwärzler, Eduard Gröller, and Harald Piringer. Litevis: Integrated visualization for simulation-based decision support in lighting design. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):290–299, 2015.
- (7) Johannes Sorger, **Thomas Ortner**, Harald Piringer, Gerd Hesina, and Eduard Gröller. A taxonomy of integration techniques for spatial and non-spatial visualizations. In *Proceedings of the Eurographics Symposium on Vision, Modeling and Visualization*, pages 57–64, 2015.
- (8) **Thomas Ortner**, Gerhard Paar, Gerd Hesina, Robert Tobler, Bernhard Nauschnegg. Towards True Underground Infrastructure Surface Documentation. In *Proceedings of International Conference on Urban Planning and Regional Development in the Information Society*, 2010.