

Hermes: Guidance-enriched Visual Analytics for Economic Network Exploration

Roger A. Leite^a, Alessio Arleo^a, Johannes Sorger^b, Theresia Gschwandtner^a,
Silvia Miksch^a

^a*Vienna University of Technology - TU Wien*

^b*Complexity Science Hub Vienna*

Abstract

The economy of a country can be modeled as a complex system in which several players buy and sell goods from each other. By analyzing the investment flows, it is possible to reconstruct the supply chain for the production of most goods, whose understanding is important to analysts and public officials interested in creating and evaluating strategies for informed and strategic decision making, for instance, adjusting tax policies. Those networks of players and investments, however, tend to be complex and very dense, which leads to over-plotted visualizations that obfuscate precious information such as the dependencies between productive sectors and regions. In this paper, we propose *Hermes*, a guidance-enriched Visual Analytics environment (named after the Greek God of Commerce) for the exploration of complex economic networks, to uncover supply chains, regions' productivity, and sector-to-sector relationships. With practical knowledge regarding guidance, we designed and implemented a visual sub-graph querying approach to extract patterns from such complex investment graphs obtained from real-world data.

We present a three-fold evaluation of the system: we perform a qualitative evaluation of our approach with three domain experts, a separate assessment of the proposed guidance features with an expert researcher in this field, and a

*Corresponding author

¹Email addresses - Roger Leite, Alessio Arleo, Theresia Gschwandtner, Silvia Miksch: `firstname.lastname@tuwien.ac.at`. Johannes Sorger: `sorger@csh.ac.at`.

²ORCID(s): 0000-0003-4427-5703

case study of *Hermes* using a bank account network dataset to demonstrate the generalizability of our approach.

Keywords: Data Visualization, Economics, Network Exploration, Supply Chain

1. Introduction

Understanding the behavior of the players in an economic network is a challenging and profitable problem. The obtained knowledge could be used to maintain a sustainable environment to support the country’s growth. For example, understanding how in a country productive sectors and regions are related allows the government to plan strategic policies that support positive economic trends. Typical tasks involve looking for weak/strong connections, supply dependencies, or growing and shrinking investment trends. While Artificial Intelligence (AI) is very useful to detect known patterns as well as expected and non-expected outliers, when it comes to reasoning and exploration of a data set, a Visual Analytics (VA) tool works as a brush to an artist. By coupling computational power with the power of the human visual system, we enable domain experts to derive insights which would not be possible with other approaches.

In this work, we focus on the problem of exploring country-wide complex financial networks, aiming at unveiling investment networks, supply chains, and economical links between productive sectors and geographical regions. We present *Hermes*, a VA exploration environment for the exploration of economic networks featuring **time-oriented** and **multivariate** data. Our system is composed of an “Overview” board to provide a high level analysis of the available data and a “Network” board for details on demand. In the latter we propose a visual approach for the definition of patterns to be searched within the network that can be defined by creating and combining “profiles”, that are groups of nodes that comply to a set of user-defined constraints. We also include guidance-enriched [1, 2] components to aid the users in the definition of queries and profiles, to enable a fully interactive visual approach to profile creation,

pattern drawing, and pattern matching. Therefore, our main contributions are:

- We designed a **visual support for the definition of designated profiles**, which can summarize multiple node types and facilitate a compact construction of query subgraphs.
- 30 • We integrated a **guidance-enriched visual subgraph querying approach** that aid the analyst in building subgraph queries from the overabundance of possible node and edge combinations.
- We conducted a **three-fold evaluation**: (1) a task-based qualitative user study with three domain experts collecting insights and guidance update rate measurements, (2) the assessment of our guidance design by an expert
35 in guidance methods for VA, and (3) a use case description demonstrating the generalizability and applicability of *Hermes* to a second domain.

2. Related Work

Hermes can be framed as a VA approach for the analysis of economic net-
40 works using a visual language to perform graph queries with guidance elements. Therefore, we describe the related work by grouping in four categories.

VA for Financial Domain: Visualization proved to be prolific in the financial domain. FinanceVis [3] presents a survey involving over 50 research articles, totaling over 70 different techniques, related to the presentation and visualiza-
45 tion of financial data. It also features an online collection of the surveyed papers for easy access. The paper concludes that visualization could empower analysts in the financial domain, and encourages further collaboration. Ko et al. [4] focused on gathering expert feedback and opinions on existing VA approaches in finance. Interviews were also targeted at understanding possible improvements
50 and preferences about visual abstractions, interaction methods, automated techniques, and data sources. When designing *Hermes*, we drew inspiration from the several visualization approaches in the financial domain [3], and from the strong motivation and appeal for more real-world based solutions [4].

Network Analysis in Finance: To enable the VA process on financial
55 networks, graph drawing and visualization played a major role, with node-link
representations being often used as an intuitive way to present such networks to
the users. Allen and Babus [5] argue that node-link graph diagrams are a nat-
ural solution for the representation of financial systems and efficiently explain
certain economic phenomena such as resources exchange. Tekušová et al [6]
60 visually supports the users in understanding complex shareholding networks. It
incorporates different economic metrics and supports the identification of ulti-
mate shareholders. VisFAN [7] deals with the observation financial transactions
between entities to identify money laundering patterns and potential frauds.
The system also allows to cluster and aggregate elements for bottom-up explo-
65 ration with both automatic and manual clustering techniques. Huang et al. [8]
present an event monitoring approach for stock markets security. It features a
linked-views approach with node-link diagrams for network analysis and a 3D
treemap for market performance visualization. The main goal of this approach
is to identify the attackers (the sources of the fraud), and further attack plans,
70 while reducing the amount of false-positives generated by AI techniques.

Origraph [9] supports network data creation, reshaping, and filtering. The
tool allows the user to execute these operations with little to no programming
according to the level of details required by the task. However, it does not
support aggregated filters such as the “profiles” presented by *Hermes*. Fuji-
75 wara et al. [10] presented a solution to automatically compose a network visual
summary based on a rank model (V2D2) and a reduction algorithm. The rank
model orders the importance of interactions and provides dependencies among
them. Even lacking user studies, the proposal automatic approach proved to be
efficient for preparing large networks for analysis. Focused on non-expert users,
80 Elzen et al. [11] navigates from network detail to overview by network elements
selection and aggregation. As a limitation, they claim to present visual clut-
ter issues with large selections (above 10) due to its clustering or community
detection algorithms.

Graphical Queries on Graphs: A typical task in the analysis of financial

85 network is to search for pre-defined patterns and subgraphs, to highlight some
transfers and behaviors of interest (e.g., money laundering, tax evasion, supply
chain management). To allow users to easily define these patterns, research
on VA for finance branched into systems that incorporated visual languages
allowing the definition of patterns and profiles of interest, which are used to
90 query the original, complex network. An example of this approach is provided
by Didimo et al. [12]. By using an already established visual language for pat-
tern definition [13], they design a system for highlighting malicious transactions
between taxpayers (that includes both individuals and legal entities). Differ-
ently from our approach, the system is designed to operate in an application
95 domain, while we propose *Hermes* as a more flexible approach more oriented
on investments rather than individual transactions. Koenig et al. [14] discuss
about the visualization of graph query search results in attributed graphs. This
means that each node belongs to a specific category derived by its attributes
(e.g., movies, actors, directors). Similarly to our approach, the users define
100 the patterns by combining the different node categories (represented as nodes)
together. To display the results, the obtained nodes are mapped to the user-
defined patterns, which helps users to keep their mental maps. On the other
hand, this approach limits the number of data items that can be shown in the
node-link diagram of query results. The problem of designing a suitable vi-
105 sual query language for graph datasets was also discussed by Chau et al. in
their GRAPHITE [15] system and, in more depth, by Pienta et al., with their
VISAGE [16]. Other than providing a platform and visualization for querying
graphs, the paper introduces a graph “auto-complete” feature to aid the user
to construct and refine queries. This inspired us to include in *Hermes* specific
110 features to assist the users in building the queries. To do so, we focused on the
use of “Guidance”, a relatively modern field of research that has been recently
characterized for VA [1].

OntoVis [17] allows the analysis of social networks by semantic and structural
abstraction. It aims at understanding who are the main actors and what are the
115 common behaviors of a social network. Structural abstraction is achieved by

the topological information such as node degree and connectivity while semantic abstraction is achieved by filtering nodes that belong to the same type. Different from the “profiles” presented by *Hermes*, that allows for fine-tuned group filters involving different numerical and categorical attributes, the semantic abstraction presented by OntoVis is limited by the existing node types.

Guidance: Guidance has roots in the field of Human Computer Interaction (HCI) [18], but Ceneda et al. [1] extended van Wijk’s model of visualization to integrate guidance into the VA process. They use existing guidance approaches from the literature to illustrate the proposed model. Ceneda et al. [19, 20] designed and implemented a data-driven guidance tool to demonstrate that guidance techniques have the potential to improve data analysis. Most recently, Ceneda et al. [2] conducted an elaborated literature research to provide an overview of how guidance is tackled in different approaches. In this work, the authors highlighted the importance of knowing the actual state of the analysis. We designed *Hermes*’ guidance features with the same “context-awareness” principals.

3. *Hermes* Design Context

The main goal of this section is to present the design context of *Hermes*, as well as to describe design and implementation choices.

3.1. *Data, User, and Tasks*

We designed and developed our approach based on the design triangle methodology [21] and Munzner’s nested model [22], by characterizing the problem at hand according to the underlying data, the involved users, and the tasks they have to carry out.

Data. The economic data that our target users analyze can be separated into two categories: the first one describes aggregate information, e.g., on national level, that report the monetary in & out flows between industry sectors within a fiscal year (“IO Tables” [23]). A transaction in such a table describes,

for example, the amount of investments (in millions of Euros) that went from
145 the “Agriculture” sector to the “Manufacturing” sector. Sectors are encoded
according to the ÖNACE specification [24]. The second data category complements
the aggregate information of IO Tables with more granular/regional
details, e.g., on the level of individual firms within a region of a country. For
this category we used the “Sabina” dataset [25]: it contains the balance sheets
150 of all the firms in Austria, including cash flows, locations, branches, and sectors.
The combination of both datasets (see Sec. 4.1) creates the **Exchange Supply
Chain (ESC)** dataset which was used during the design and implementation
of *Hermes*. The dataset features 172 nodes and 20,709 edges.

Users. *Hermes* target users are financial and economical experts involved in
155 the development of regional policies and investment plans, for instance in order
to economically revitalize a region. Users are expected to be proficient in sta-
tistical evaluation of tabular data, but are not expected to have any experience
in visual analysis or network analysis.

Tasks. The task that *Hermes* aims at supporting, is to improve the under-
160 standing of national economic investment networks based on the integration of
aggregated and granular data.

We identified four sub-tasks: the exploration of the economic flow network
existing between sectors (**T1**) and regions (**T2**), the evaluation of how sectors
relate and impact each region (**T3**), and supply chain exploration (**T4**) (e.g.,
165 predominantly buying/selling regions).

3.2. Research Challenges

In the context of *Hermes*, we are dealing with several challenges that result
from the underlying data as well as the tasks that should be enabled on them.

Dense Dynamic Networks: Large economic networks are dense and dy-
170 namic, with large and frequent fluctuations. While automatic approaches can
still be used to detect outliers and confirm known patterns, VA techniques sup-
port economic networks analysis, as they enable users to interactively drill down
on the dense data with multiple views, improving the exploration process.

Multi-scale Context: the analysis and exploration of transactions on and
175 across different scales, i.e., regional, national, global, requires visual and ana-
lytical scalable solutions.

Geo-political Context: To better understand transactions, we need to
consider the motivations that guide them. It is known that geopolitical, geo-
graphical, and economical contexts influence this behaviour, e.g., buying from a
180 local or a distant supplier, or the presence of a quarry or oil field. Considering
the ever changing local and global scenarios, an analysis tool must be flexible
in conveying these different contexts.

Temporal Trends: Economic network analysis involves the observation of
stable behaviour over time (e.g., distinguishing growing regional sectors from
185 shrinking and stable ones). However, there are many aspects (scale, scope,
arrangement, viewpoint) concerning time-oriented analysis that need to be an-
alyzed efficiently due to their complex data features [26].

3.3. Requirements

By combining our tasks (see Section 3.1), the research challenges (see Sec-
190 tion 3.2), and with the collaboration of three economists, we derived four main
requirements for our system. The economists that collaborated during the re-
quirements and the design development were not participants of the evaluation
presented in Section 5.

R1: Overview of region and sector distributions. An economic
195 environment is in constant evolution: the transactions between its major players
(the companies) can leave a significant and lasting impact on the economy of
the region they belong to as well as their foundation and demise. An overview
of these aspects is important to understand the main sectors of an economy and
supports insights such as the most and least prominent sectors among different
200 regions.

R2: Multi-Scale network exploration. Economists have to understand
the national investment and supply chain network to better argue about possible
investments, propose tax remodulations, and other types of economic policies.

The system must enable network exploration on the different levels of detail on
205 which the information is available. Moreover, the system must implement filter
techniques to support the exploration of such a dense network.

R3: Managing queries. During network exploration, analysts regard
different aspects of the data by adjusting various filter settings. Often it is nec-
essary to compare and cross-reference the results of such queries. A competent
210 system must thus support the creation, visualization, and comparison of such
query profiles.

R4: Identification and tracking of monetary flows. Understanding
the monetary flows between sectors and regions could give a significant advan-
tage in many scenarios, including investment planning, supply chain exploration,
215 and fraud detection. The system should support the users in the identification
and tracking of the flows in the transactions network, uncovering paths and
patterns pertinent to the application domain.

4. *Hermes* Design and Implementation

We designed several prototypes that we used to explore the design space in
220 terms of visualization approaches and interactions methods. During different
phases of the project development, the prototypes were inspected and evalu-
ated by our three economists collaborators that supported the requirements
definition (see Section 3.3). Our approach was developed using web technology
(HTML, CSS, Javascript), as well as frameworks, services, and libraries such as
225 Angular, JQuery, and D3.js. A sketch of the system architecture can be seen in
Fig. 6.

4.1. *Data Model*

Hermes data model is represented as a network, whose nodes, called **View**
Nodes, represent a productive sector of a specific region (e.g., “Agriculture” in
230 region “Vorarlberg”). Nodes also have attributes, including cash income flow,
outcome flow, and growth rate. We obtain the growth rate by calculating the

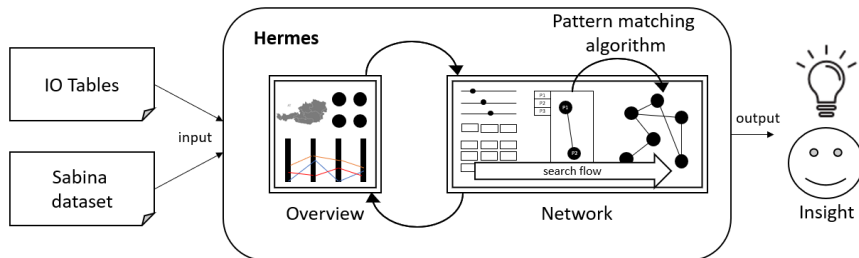


Figure 1: *Hermes* workflow. Using the two types of input data (left-hand side) *Hermes* presents two connected visualization boards (Overview and Network), which support interactive exploration and insights generations (right-hand side).

difference of the cashflow (income flow + outcome flow) between the current year (y) and the last year ($y-1$). Nodes are linked if a financial connection (e.g., investment) exists. We obtain this model from our data sources (see
 235 Sec. 3) as follows: based on the data combination described by Fritz et al. [27], we link the details about national companies (i.e., expenditures, amount of employees, sectors, locations) to the monetary income and outcome flow of each productive sector of different regions, enabling us to estimate how much money was exchanged between the different sectors and regions. We can, thus, define
 240 different *profiles* to filter these nodes. A *profile* is a set of constraints that identifies a subset of nodes. These constraints might include a group of sectors, a group of regions, in- and outflow ranges, as well as a growing rate range.

4.2. Pattern Matching Algorithm

To create graph patterns, we define and combine specific profiles. These
 245 profiles are represented as profile nodes and the user can link them using the Network Board (see Sec. 4.3.2). Once a pattern has been constructed, we can search for matching subgraphs in the data model. We define $G_t = (V_t, E_t)$ as our data model graph (i.e., the target graph) described in Sec. 4.1, with V_t the vertex set and E_t the set of ordered pairs of nodes representing the graph edges.
 250 We also define $|V_t| = n$. $G_p = (V_p, E_p)$ is the pattern the user constructs with

profile nodes (i.e., the pattern graph). We use our pattern matching algorithm to look for a subgraph $G_r = (V_r \subseteq V_t, E_r \subseteq E_t)$, i.e. the result graph. This algorithm conducts the following operations: since G_p nodes represent profiles, which identify a set of nodes in V_t , first **(OP1)** for every profile node $v_p \in V_p$ it
255 extracts the nodes $v_t \in V_t$ that match the constraints defined in v_p and assign them to V_r ; **(OP2)** for every node now in V_r , it checks whether its outgoing edges belong to E_r . An edge $e_t = (u_t, v_t) \in E_t$ belongs to E_r if $u_t, v_t \in V_r$ and they were extracted from two different profile nodes in V_p connected by an edge $e_p \in E_p$.

260 It’s immediate to show that our algorithm has a quadratic asymptotic time complexity. *OP1* has to be repeated for each node in V_p , therefore has a complexity of $O(|V_p|)$. We can perform *OP2* in $O(d)$ time for each node, with d the out-degree of a node (i.e., the number of outgoing edges). Therefore, considering that from V_p we extract the nodes to insert in V_r , the time complexity of the
265 algorithm is $O(|V_r|d)$: $|V_r|$ is at most n , when the queried subgraph includes all the nodes in V_t , and d is at most $n - 1$, when v_r shares an edge with all the other vertices in the graph. Therefore, the asymptotic complexity is $O(n^2)$. In practice, we could find that also with very dense graphs our implementation was able to provide the requested subgraph in very short time, thus enabling
270 an interactive exploration (see Sec. 6). We would like to remark that even if the general pattern matching problem is NP-hard, in our case the problem is greatly simplified because the mapping between the profile nodes and the nodes in the target graph is already known, since is defined by the user.

4.3. Different Boards for Different Tasks

275 In this section we describe our design of two combinations of different visualizations (i.e., “boards”) serving different tasks. In both boards, we used “Color Brewer” [28] to support distinct color pallets for representing different regions. The “multiple boards” approach supports “R2: Multi-Scale network exploration”.

280 4.3.1. Overview Board

We designed a fully connected multiple coordinated view board (see Fig. 2) to provide a good overview of region and sector distributions (R1) and to support business contextualization. Views are connected with brushing and linking and filter operations in one view are reflected in all the other views accordingly. Besides providing a good overview, this board is designed to support reasoning about how regions are composed of sectors and how sectors are represented in the different regions. Next, we explain our design choices and the possible interactions for each of the views.

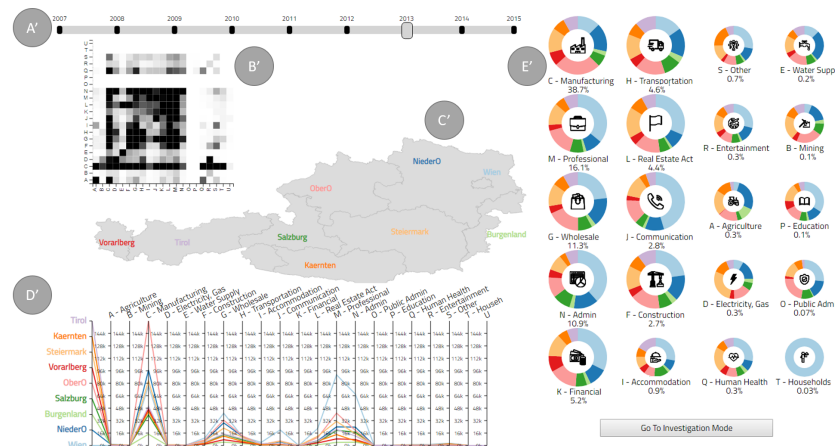


Figure 2: Screenshot of *Hermes*: Overview Board. (A') Time Slider: allows the selection of different years. (B') Connection Matrix: gives first impressions about how intense the sectors are connected on a national level. (C') Regions Map: shows the color code and the geographical location of the different national regions. (D') Parallel Coordinates: support users to understand how sectors are distributed among regions. (E') Sectors Donut Charts: shows how regions are distributed among sectors.

A': Time Slider. The main goal of the time slider is to select data from different years for analysis. The user can drag and drop the slider on any year to load the data of this year.

B': Connection Matrix. We show a preview of connections between sectors for the whole country by using an adjacency matrix. We choose adjacency matrices to visualize a dense network [29]. Adjacency matrices also have the ad-

295 vantage that they emphasize an additional variable (i.e., in our case the amount
of money transacted between two sectors) instead of connections and paths as
node-link diagrams do. Moreover when compared to node-link diagrams, ad-
jacency matrices make better use of the space available and they present less
visual clutter. Each square represents a connection and the shade of the square
300 encodes the amount of money transacted. Dark areas in the matrix can be spot-
ted immediately which represent big quantities of money transacted. We also
implemented a mouse over feature that displays who is the source and target for
each matrix connection square, to make precise information about connections
between nodes [30] and amount of money involved available.

305 **C': Regions Map.** The Austria national map including different regions
not only works interactively by selecting different region's selectors directly in
the map, but also supports the contextualization of physically neighboring re-
gions. Therefore, it also supports presentation purposes.

D': Parallel Coordinates. Each colored line represents a different region
310 while the axes represent the amount of companies of each sector. To support
different overview tasks, it is also possible to represent other features in the axes
such as the total revenue. With this view it is possible to observe how regions are
composed by sectors and how sectors are distributed between regions. We opt
for using parallel coordinates due to the fact that it supports the representation
315 of multiple dimensions. The axes share a common scale to support comparison.

E': Sectors Donut Charts. A donut chart is essentially a pie chart with-
out the center area. We used the central area for plotting the symbol of each
sector. Below each chart, we plot the title of the sector and the percentage that
represents the portion of companies within this sector with respect to the all
320 companies in the country (in the selected year). While donut charts should not
be used for precise comparisons, we chose them for two reasons (1) to create
an area dedicated to easy sector overview and selection (comparable with what
view Fig. 2: C') does for regions) and (2) to get an immediate impression of
the main regions a sector is located in. Also, their aspect ratio allows for the
325 representation of multiple sectors at compact space. In this view, we distinguish

between two sizes of donut charts, smaller charts represent sectors that had less than one percent of the national cash flow while bigger charts represents sectors with more than that. This feature facilitates the task of identifying the main sectors of the country. It is possible to select values (sectors and/or regions) from the Overview Board to filter the data that is available in the Network Board, which implements the overview + details on demand mantra [31].

4.3.2. Network Board

During this subsection we refer to view labels from the Fig. 3. The Network Board is divided into views A, B, and C to support a three step workflow. In view A, the user creates profiles of sectors and/or regions with specific income and outcome properties as well as growing rates. For example, a profile could be “sectors from Burgenland with a cash income of more than 30 and less than 100 million Euros”. In view B, these profiles can then be used as network nodes to construct a network pattern to query the data. The result of a query is displayed in view C. Views D and E present different exploration features that support view C. The workflow presented by the Network Board aims to support “R3: Managing queries” and “R4: Identification and tracking of monetary flows”.

A: Profile Creation View. The profile creation view is divided into two parts: (1) the guidance-enriched sliders to filter income, outcome, and growing rates, and (2) the guidance-enriched tiles to select regions and sectors. Inspired by the Scented Widget [32], the guidance-enriched sliders are designed to increase user familiarity with the data during the profile creation process. Three sliders are presented: income cash flow, outcome cash flow, and growth rate. The income and outcome cash flow dimensions filter how much money a possible sector or region handles while the growing rate filters the growing (or shrinking) rates of sectors and regions. Growing rates are derived by comparing the analyzed year (y) to the previous year ($y-1$). The guidance aspects to facilitate informed filtering are described in Sec. 4.4. Defined profiles are available in view (B) and can be selected as a node to construct a network pattern. Unique profile IDs are assigned which contain a unique letter + the amount of represented

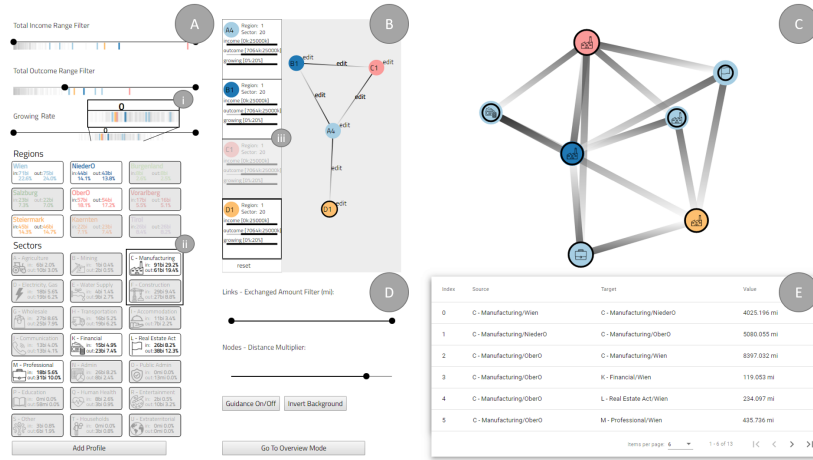


Figure 3: Screenshot of *Hermes*: Network Board. (A) Profile Creation View: Profiles created in this view are used to query network patterns in view (B). View (A) is divided into two parts: (i) guidance-enriched sliders and (ii) region/sector filter selection. (B) Guidance-Enriched Pattern Search: here the user can link together profiles to create patterns to be matched. The query result is displayed in (C). (D) Control Panel: this panel allows users to regulate some features from view (C). (E) Dynamic Table: this view shows raw transaction data.

regions/sectors. Profile definitions might be very broad as well as very specific. For instance, a profile which filters only the income dimension to more than 12 billion Euros will match three nodes in the data set described in Sec. 3.1. On the other hand, if we filter to less than 12 billion Euros, 160 nodes would match the profile.

360

B: Pattern Search View. In the Guidance-Enriched Pattern Search View, we have (1) a list of created profiles on the left and (2) a drawing canvas on the right. This view allows users to query for network patterns by selecting profile nodes in (1) and drawing connections between them in (2). A visual summary of the created profiles is shown in (1): a circle filled by the regions' color (or gray if multiple regions are involved) and the profile's ID together with small symbol representing the intervals of income, outcome, and growing rate for this profile. The drawing canvas (2) allows the user visually to recognize the query patterns efficiently [33]. It supports various profile connection combinations.

365

370 The canvas also allows for interactively deleting and adding profile nodes and
edges as well as changing the direction of edges. Edges might be uni-directional
or bi-directional. The generation of profile nodes allows the user to create a
compact representation of complex constraints, i.e., multiple sectors and regions
with specific characteristics can be combined into one node, which facilitates
375 the construction of network structures to query for and improves clarity of
these structures. We present the combination of profiles in (2) as a node-link
diagram not only to preserve visual analogy between the constructed pattern and
the network data represented in (C), but also to better support path-findings
tasks [34]. Moreover, compared to other visualization of network data, such as
380 adjacency matrices, node-link diagrams are better suited to communicate and
memorize network topology and connectivity [34]. These tasks are essential for
analyzing economic flow.

C: Node-Link View. We opted for a node-link view instead of adjacency
matrix due to two main design advantages: (1) it supports visual matching
385 between the pattern drawn by the user in view B and the query results and (2)
it supports better reasoning about indirect paths between two nodes (R2 and
R4) [30]. The compact network pattern created in view (B) is used to query the
data set for sub graphs that match this pattern. The results are then shown in
the Node-Link View. However, this view shows a faithful representation of the
390 data set without collapsing multiple sectors into one node, as it is done in view
(B) to facilitate pattern construction. Thus, each node in this view represents
a specific sector of a specific region and edges represent the yearly money flow
between them. The size of nodes is proportional to the amount of money they
handle (income & outcome) while the ring close to the outer border of the node,
395 indicates the grow rate by representing the size of the node in the previous
year (see Fig. 4). In other words, if the ring of a node is bigger than the node
itself, the amount of money handled by this sector was decreasing since last
year. On the other hand, if the ring is smaller than the node, this sector has
a growing trend (see Fig. 4). The amount of money exchanged between sectors
400 is encoded by edge width. To define this width, we again use a logarithmic

scale to avoid extreme visual outliers. Following Holen and van Wijk’s [35] recommendations, we avoided standard arrow representation and, instead, chose a dark-to-light encoding. The money flows from the node at the dark side of the edge to the node at the light side. For bi-direction edges both sides are dark (see leftmost node in Figure 3). All nodes and edges can be selected for viewing detail information in the Dynamic Table View (see view (E)).

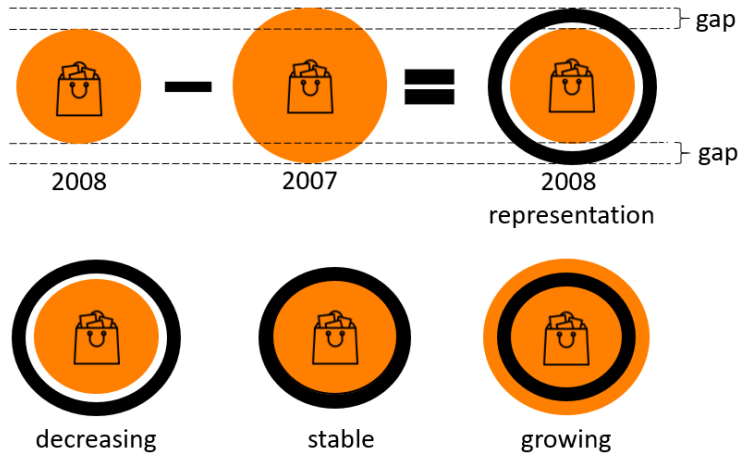


Figure 4: Visual encoding of the growing rate of a sector. The first row shows the comparison operation and how it is reflected in the visual encoding of a sector node. Using this metaphor, we can visually represent shrinking, stable, and growing sectors (second row).

D: Control Panel. The control panel was implemented to facilitate the inspection and refinement of the network presented in view (C). The first slider allows filtering edges by amount of money and the second slider allows re-scaling the network by increasing or decreasing the forces between the nodes. Since edges directions are represented by a dark-to-light gradient, we provide the option to invert the the background of the node-link diagram to black, which facilitates the identification of receiver patterns.

E: Dynamic Table. During our problem analysis together with experts, we learned that they prefer to rely on detailed information presented in a spreadsheet-like representation. While tables are not suited to provide a good overview of the network structure, they are good in presenting all the detail

information of specific sectors and regions. Thus, we implemented a dynamic table to provide experts with these details on demand. The dynamic table is
420 populated with information about all edges visible in view (C). Selecting edges in this table highlights their representations in view (C).

4.4. Guidance Aspects

In this subsection we describe the different degrees and features of guidance that are provided by *Hermes*. Guidance by itself is a computer-assisted process
425 that continuously reduces the knowledge gap that may block the data exploration and analysis. It supports the users to make sense of the data on their own [1, 19, 2].

A huge number of profile combinations and constructed network patterns would be feasible, but only a small set of them make sense for the data at
430 hand. To avoid frustrating trial and error interaction until a meaningful pattern for querying the data set could be found, we needed to design *Hermes* with features that guide the user in this pattern construction. *Hermes* includes four different features aimed at providing guidance to the user, and thus, facilitating a meaningful and informed usage of the tool. The following indexes “i,ii,iii”
435 refer to Fig. 3 and “iv” refers to Fig. 5.

i: Slider Ticks. Sliders in the Profile Creation View are augmented with ticks colored with respect to the different regions (see Fig. 3: i). One tick represents one sector in one region. Hovering these ticks makes a tool tip window appear which shows full textual information about the sector, region, income,
440 outcome, and growing rate. Ticks are grayed out when they do not match all defined filters. Thus, this visualization of ticks at the point of interaction (i.e., the slider) guides the user in making a meaningful selection of sectors for profile construction, by a priori avoiding constructions that would not lead to any results as these constructs cannot be found in the particular data set. This
445 kind of guidance (i.e., adding information that facilitates the interaction and selection of the data) meets the guidance degree of orienting [1] which is aimed at building and preserving a user’s mental map [36].

ii: Selection Tiles. The tiles for selecting regions and sectors work in accordance with the slider filters, for including or excluding regions and/or sectors from the filtered selection. Options that, in combination with the other filter settings, do not exist in the data are grayed out (see Fig. 3: ii). This, again, helps to avoid the construction of non-existing profiles. Since this feature limits the users interaction possibilities (i.e., grayed-out tiles cannot be selected), it corresponds to the guidance degree of directing, which usually offer the user a pre-selected pool of alternative next steps for a given task [1].

iii: Guided Network Pattern Construction. While the two previous guidance features help the user to define meaningful profiles, the combination of these profile nodes into a network pattern that can be found in the data, is also a laborious task if performed without guidance. Thus, the Pattern Search View (see Section 4.3.2 (B)) does not allow for the construction of non-existing patterns. This is achieved in a two-step process: First, every time the user selects a profile node that is part of the constructed pattern in the drawing canvas (in Fig. 3: B) in order to add a connection to another profile node, the system runs a query that looks for possible connections from this node in the data set, and disables the selection of all profile nodes that would lead to empty query results. For example, a given “profile C” that would lead to empty query results when connected with “profile A”. However, it might be possible to connect it with “profile B”. In Fig. 3: iii, the node D1 is selected in the pattern drawing canvas and thus, profile C1 is disabled in the list of available profiles by the system. Again, this feature limits users in their possibilities to construct arbitrary patterns, which makes it directing guidance [1].

iv: Switching to Bipartite Layout. Another design choice that helps orienting the user [1], is that the network presentation in the Node-Link View switches to a bipartite layout (see Fig. 5) every time the user constructs a profile combination in the Pattern Search View that connects only two profiles. The automatic adaption of the layout creates an additional awareness in the user that he/she is now investigating a bipartite graph. While this is obvious from the query pattern constructed in the pattern canvas, it can be hard to spot in




Figure 5: In this picture, we see a bipartite layout result of two profiles. On the left-hand side, the first profile is representing the manufacturing of each region and, on the right-hand side, the second profile represents the transportation sector of each region. Based on this result, we can investigate how much each manufacturing region spends with each transportation region.

the Node-Link View which nodes from the query result belong to which profile.

480 Moreover, the layout facilitates the flow analysis between the two groups.

5. Evaluation

To assess and evaluate the effectiveness and limitations of *Hermes*, we conducted a qualitative user study, in the form of an interview with three domain experts (P1, P2, P3) [37, 38]. We first describe the user study design (see
 485 Sec. 5.1) followed by the participant background (see Sec. 5.2); afterwards we discuss the tasks and the immediate findings of our study. To complement our

study, we interviewed a research expert in guidance (P4), who discussed his/her thoughts about the theoretical and practical implications of the guidance features of *Hermes* (see Section 5.5). Furthermore, we also did examination of the transferability of *Hermes* to Financial Fraud Analysis (FFA) with a domain expert (P5).

5.1. Study Design

We conducted our user study as a task-based evaluation: the objective was to evaluate and observe how users operate *Hermes* when asked to solve various tasks. Our study design consists of individual interviews followed by a hands-on session to be carried out by up to two participants at a time. Each interview is structured as follows: (1) participants background interview, (2) introduction of overview and Network Board, (3) hands-on testing, (4) task-based evaluation, and (5) final general feedback discussion. Each session lasted about 90 minutes. We aimed at making the users familiar with *Hermes* before asking them to perform the task-based evaluation individually and (mostly) unguided. For the whole duration of the session, participants' interaction with *Hermes* was recorded. They were also asked to think-aloud in order to follow their reasoning during the evaluation. The tasks chosen are different for the overview and Network Boards and were designed to evaluate whether *Hermes* could fulfill its design requirements (see Sec. 3.3). We identified four tasks per board (see Sec.s 5.3 and 5.4), structured according to the analytical task taxonomy by Andrienko and Andrienko [39], which highlights differences between elementary and synoptic tasks.

Furthermore, to support insight analysis, we categorize the obtained user insights by following Klein's model of intelligence analysis [40]. Our taxonomy includes: **Connection** insights, obtained by combining data from two or more views; **Coincidence** insights, which include relationships between elements or events that are obscured at first, but that the users can perceive as connected; **Curiosity** insights happen when users do not find information when they expect to, and this pushes them to find a reason why. Finally, in order to fully com-

prehend the impact of our guidance techniques in our study, we implemented in our prototype a counter that would keep track of how many times the system updated the interface to guide the users (see Sec. 4.4). We called this measurement **Guidance Updates**. The number of insights and guidance updates are displayed in Fig. 6 for each task. During the evaluation session with P1, P2, P3, and P4, we used *Hermes* with the *ESC* dataset (see Sec. 4.1). The FFA dataset was exceptionally presented to P5, who is domain expert in the field (see Sec. 5.6).

5.2. Participants' Background

P1 and P2: The participating experts (not involved in the design) can all be considered potential *Hermes* target users (see Sec. 3.1). The two participants are employees from the economic and trade policy department of the “Austrian Federal Economic Chamber” (*Wirtschaftskammer Österreich* or WKO). They analyze the trends of business interests in Austria as well as evaluating and predicting the effects economic policies and regulations throughout the different regions of the country. Both participants use Microsoft Excel [41] as well as “Python” and “R” to explore and present their data. However, they report a lack of visual tools that support their work. Therefore, it is expected that they also lack visual guidance. One participant was familiar with “Tableau”, however does not apply it in his work context.

P3: Our third participant is an economist, who works as the global treasury controller within an international company. In contrast to the other two participants, who looked after the productivity of entire regions and multiple companies, P3 is interested in the company’s business only. Therefore, P3 is looking for a system that could support investment decisions with information about product supply chain networks. The most used tool by P3 is Microsoft Excel Sheets [41] and P3 claims to be familiar with simple visual data representations existing o it.

P4: Aiming for a focused qualitative evaluation of the guidance aspects in *Hermes*, we invited a VA expert researcher with a research focus on guidance.

The participant was used to develop VA techniques.

P5: Our fifth participant is a financial fraud expert that works for a national bank. P5 claim being familiar with a good range of visualization techniques and, on her/his work, s/he uses visualizations for presentation and exploration
550 proposes.

5.3. Evaluation of the Overview Board (P1, P2, P3)

We designed the Overview Tasks (**OT**) to evaluate the first of the two *Hermes*'s boards, which was designed to fulfill requirement R1 (see Sec. 3.3).

555 The tasks are the following:

OT1: Which are the biggest sectors in Austria?

OT2: Which regions present more manufacturing companies?

OT3: What is the main region? Why?

OT4: By looking and interacting to the Matrix View (Fig. 2: B'), what
560 type of conclusions could you take?

All the participants concluded the tasks before the specified time, and reported correct answers. Connection insights were the most common insight category during the evaluation of the Overview Board. However, OT4 did not present any connection insight because it is an one view task. Overall, we interpret this result as a positive consequence of presenting a fully interactive
565 multiple coordinate view approach that shows different dimensions of the same data elements. The combination of different selection and filtering features was the main aspect explored during the performance of OT1, OT2, and OT3.

OT3 was especially designed to motivate value thinking by presenting the words “main region” on its specification without describing what was the comparison context encapsulated by the word “main”. The participants could define
570 different meanings to the word “main”. Still the three participants could explore their hypothesis and reasoning about the different possible outcomes of this interpretation, thus, revealing the flexibility of the tool.

575 5.4. Evaluation of the Network Board (P1, P2, P3)

For this board we also conceived four tasks. To evaluate R3 (see Sec. 3.3) we designed two similar network tasks (NT1 and NT2). The goal is understanding the impact of guidance during the profile construction phase. To do so, we deactivated the guidance features in *Hermes* when evaluating NT1 and reactivated
580 them on NT2. In addition, we asked our participants to tackle two open tasks (NT3 and NT4), which allowed for multiple correct answers. The goal is to observe how many insights and conclusions the participants could draw within the given 5 minute time frame. NT3 and NT4 aim at supporting R2 and R4 analysis (see Sec. 3.3).

585 **NT1:** Excluding the region “Oberösterreich”, which View Nodes are the sectors with the four highest *income*? How are they connected? The task has to be completed without guidance support.

NT2: Excluding the region “Oberösterreich”, which View Nodes are the sectors with the four highest *outcome*? How are they connected? In this task
590 the guidance features are enabled.

NT3: Choose two regions and explore their strongest connections. Discuss your findings.

NT4: Choose a region and search for growing sectors and shrinking sectors. Discuss your findings.

595 Without guidance support, two out of the three experts were not able to finish NT1. P2 achieved it but past the allotted 5 minutes mark. Without visual guidance during profile creation, participants were not able to precisely set the income slider to a convenient income range, that led to the inclusion of a large amount of nodes (noise) in each of the created profile nodes. As
600 a consequence, the resulting matched subgraphs did not help comprehending the data and solving the task. On the other hand, the attempts at narrowing the income range failed at including the minimum amount of searched nodes specified by NT1.

With guidance support (NT2), all participants were able to finish the similar
605 task within less than two minutes. Guidance thereby facilitates the determi-

nation of a threshold that incorporates data points relevant to the task. By observing the colored ticks under the outcome slider, users could understand how many nodes they were selecting, thus setting a “convenient” slider range. Moreover, since the excluded region is color coded on the Slider Ticks, the participants could exclude samples from the selection without necessarily interacting with selection region filters (see Sec. 4.3.2), thus speeding up the profile creation process. Tasks NT3 and NT4 are open exploration tasks that were designed to motivate open insights. Both tasks resulted in the most significant amount of insights across all tasks (see Fig. 6).

Overall, the participants provided positive feedback about *Hermes* and complemented different aspects: “*The guidance features are really useful*” (P1), “... compared to the technology I use, to perform network exploration tasks, this approach is much faster and appealing” (P3), or “...to perform the same tasks with our current approaches is almost impossible” (P2). P3 also commented positively about the possibility to switch to a bipartite layout: “...it is good when I search for less groups because the layout gets more organized”. All participants gave positive feedback on the temporal aspect derived by the growing rate slider. It presented a main role during many of the participants’ insights elaboration.

Considering the feedback and the results of the task evaluation, the data appears to support our claims, confirming that *Hermes* fulfills the design requirements. NT1/2 show the capabilities of *Hermes* when handling queries, with the guidance support (R3). The amount of findings in NT3/4 and the feedback gathered suggest that *Hermes* succeeds in enabling multi-scale network exploration (R2) and the tracking of monetary flows and region/sector connection (R4).

5.5. Guidance Expert Feedback (P4)

The focus of this evaluation was to have P4 to interact with the tool to collect expert feedback about the guidance aspects and not measure insights or answers precision. After performing the complete evaluation, during the

post tasks interview, the participant mentioned several positive points. S/He complemented different aspects such as the short interaction response time of the Overview Board and the tasks enabled by the Network Board. S/He highlighted the importance of the guidance features during the profile creation (see Sec. 4.3.2 view A). Starting with the Slider Ticks (see Fig. 3: i) s/he commented: “...with the guidance it works perfectly to select the profile that you want to see, otherwise it is complicated to find what you want. It gives you hints about the data that you are selecting”. Concerning “Selection Tiles” (see Fig. 3: ii) and “Guided Network Pattern Construction” (see Fig. 3: iii), P4 considered this guidance aspect as something in between the *orienting degree* and the *directing degree*, more inclined to directing degree (see Sec. 4.4). Concerning the “Switching to Bipartite Layout” (see Fig. 5: iv), s/he commented: “it is a good choice that the system supports this automatic decision. It is visually clear for me that this is a better layout for comparison tasks and group understanding”.

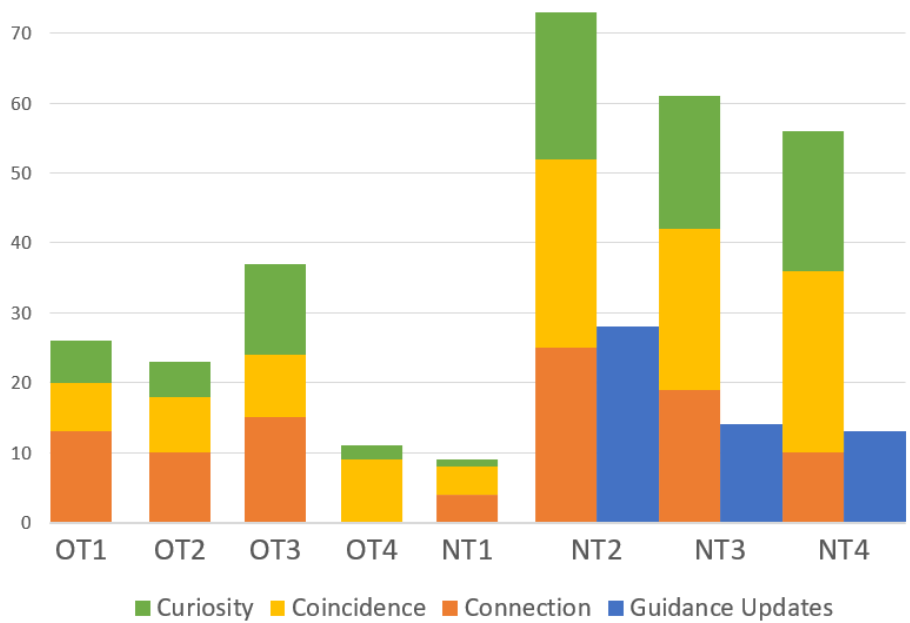


Figure 6: The overall number of and guidance updates we identified over the different board tasks. We report both the exact numbers in a table and an area chart to easily extract trends from it.

650 When asked about future recommendation for guidance features, a point raised by our participant was the limits of guidance (see Sec. 6). How much guidance is enough? S/He commented that guidance should always serve a propose, a task. However, when trying to add guidance to many tasks in the same tool, guidance features trend to overcome each other, leading users to
655 confusion. Based on that comment, s/he mentioned that all proposed tasks that were covered by guidance aspects, were not disturbing each other and, s/he would not suggest any further guidance implementation, unless the system is extended to support new tasks. In the same sense, s/he also complemented that the profile creation view does not enforce a specific workflow, avoiding
660 influencing users to limit their profile creation: “...it is good to avoid steps repetition during exploration tasks”.

Overall, s/he was surprised to see different degrees of guidance in *Hermes*. Her/His favorite features were the guidance implemented in the profile creation and the pattern search view. S/He highlighted several times how much easier
665 and better it was to have guidance features supporting her/him during these tasks.

5.6. Applying *Hermes* to Bank Account Networks (P5)

We demonstrate the *generalizability* of our approach by applying it to a different domain, i.e., the analysis of financial fraud. The FFA data is a real-world
670 dataset for which we render bank accounts as nodes and money transaction between accounts as edges between them. Transaction are ranked according to a metric that measures its how likely it is to be part of a fraudulent behaviour calculated by a bank internal automatic algorithm. *Hermes* could be used by fraud investigators for tasks such as identifying terrorists, bank credit exploration, or
675 fraud analysis. For these cases, profiles of bank accounts could be composed of features like age, balance, transaction frequency, account growing rate. With the combination of profile definition and profile-based network pattern queries the investigators would gain a supportive and flexible ally.

We also had the opportunity to present *Hermes* applied to FFA data to

680 a FFA expert (P5), and to collect positive feedback. P5 was able to interact and understand all *Hermes*' features. In particular, the expert appreciated: (1) the potential of flexible profile construction coupled with the network pattern search for detecting networks of fraudsters, (2) the quick pattern search response of *Hermes*, and (3) the guidance aspects during profile creation - which s/he
685 assessed essential for understanding what s/he was selecting and allowing for an informed profile creation.

6. Discussion and Future Work

In this section we list the limitations of *Hermes* and we discuss some general aspects about guidance that came up during the discussions with the guidance
690 expert. Moreover, we derive directions for future work.

Limited Support of Profile Membership Identification. We provide a feature that changes the layout of the node-link diagram to help the user understand which nodes belong to which profiles (see Sec. 4.4: iv). By now this feature is limited to bipartite graphs (i.e., including only two profiles). Because
695 this feature was very well received in the evaluation, we plan to incorporate a similar functionality also for queries including more than two profiles. This comes with additional challenges of how to arrange the graph layout when the number of profiles exceeds a certain limit. Radial layouts, for instance, might be implemented to address this problem.

Connections Between Profiles are AND Operations From our re-
700 quirement analysis we derived the need of analysts to search for patterns where all node types of one profile are connected to all node types of another profile. Thus, connecting two profiles in *Hermes* queries the data for network structures in which nodes of different profiles with an AND (i.e., all nodes of one profile
705 are connected to all nodes from the other profile). However, new ideas came up in the evaluation and thus, we want to further investigate if connecting two profiles with an OR operation would provide additional benefits to analysts.

Pre-defined Profiles & Network Structures. Another idea stemming

from our evaluation with domain experts is to provide pre-defined profiles as
710 well as pre-defined network structures of common queries. For example, ex-
perts suggested a profile that involves only elements with a high grow rate or a
profile of elements with a high outcome. They also suggested that pre-defined
network structures would support an even faster analysis. However, to pro-
vide pre-defined profiles and network structures, domain knowledge is required.
715 Moreover, it is not clear how such “ready to use” building blocks influence users
in their analysis and if they would limit exploration breadths [42].

Performance Considerations. As we described in Sec. 3.1, we experi-
mented with the ESC dataset, a fairly dense graph. Despite this, thanks to the
simplicity of our approach (see Sec. 4.2), we observed that *Hermes* featured low
720 response times also with more complex queries. During our expert evaluation
and internal testing, we tested queries that included stars (a central node con-
nected to some neighbors), cycles, disconnected components, and combinations
thereof up to 21 profile nodes (whose constraints identified several nodes in the
model graph). We measured the response times, from the pattern extraction to
725 the display of the matched sub-graph, which peaked at 67 milliseconds.

Classification of Guidance. It is not always unambiguous how to clas-
sify the provided guidance and even what should be considered as guidance.
However, the expert agreed that quite often features might present a mixture
of different degrees. For instance, all Selection Tiles that do not exist in the
730 data in combination with the given filter settings are grayed out. This might
be considered just a hint for orientation (i.e., orienting guidance), but it also
presents users with a list of meaningful alternatives to proceed with, which is
considered directing guidance. Sometimes basic visualization features comply
with the definition of guidance. For example, in a node-link diagram, encoding
735 the connectivity degree of a node in the node’s size to emphasize how well a node
is connected might also be considered some sort of orienting guidance. How-
ever, these basic features are common practice in the visualization community
and, therefore, are not considered to be guidance. Further research is needed to
better distinguish guidance.

740 **7. Conclusion**

We developed *Hermes*, a guidance-enriched VA approach to support national economic supply chain analysis on an overview level as well as on a detailed network analysis level. The design and implementation of *Hermes* was inspired by our own research experience and by tight collaboration with economy domain experts, with whom we elaborated a number of design requirements. Our approach consists of two connected boards which provide overview and network views augmented with various guidance aspects to support economy experts with effectively exploring and analyzing complex data.

We designed and evaluated *Hermes* with two real-world data sets and five domain experts (three economists, one FFA expert, and one guidance researcher expert), from whom we obtained positive feedback and further ideas on various aspects of our design. The added value of our approach is: (1) the analysis of insights gained by our evaluators, (2) a positive evaluation of the guidance aspects by a guidance researcher expert, and (3) the demonstration of the transferability of *Hermes* to FFA.

8. Acknowledgments

This work was partially supported by the Research Cluster “Smart Communities and Technologies (SmartCT)” at TU Wien and the Austrian Science Fund (FWF), grant P31419-N31 (Knowledge-Assisted Visual Analytics (KnoVA)).

760 **References**

- [1] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, C. Tominski, Characterizing guidance in visual analytics, *IEEE Transactions on Visualization & Computer Graphics* (1) (2017) 111–120.
- [2] D. Ceneda, T. Gschwandtner, S. Miksch, A review of guidance approaches in visual data analysis: A multifocal perspective 38 (3) (2019) 861–879.

- [3] M. Dumas, M. J. McGuffin, V. L. Lemieux, Financevis.net - a visual survey of financial data visualizations, in: Poster Abstracts of IEEE VIS 2014, 2014, poster and Extended Abstract.
URL <http://financevis.net/>
- 770 [4] S. Ko, I. Cho, S. Afzal, C. Yau, J. Chae, A. Malik, K. Beck, Y. Jang, W. Ribarsky, D. S. Ebert, A survey on visual analysis approaches for financial data, in: Computer Graphics Forum, Vol. 35, Wiley Online Library, 2016, pp. 599–617.
- [5] F. Allen, A. Babus, Networks in finance, The network challenge: strategy, profit, and risk in an interlinked world 367.
775
- [6] T. Tekušová, J. Kohlhammer, Visual analysis and exploration of complex corporate shareholder networks, in: Electronic Imaging 2008, International Society for Optics and Photonics, 2008.
- [7] W. Didimo, G. Liotta, F. Montecchiani, P. Palladino, An advanced network visualization system for financial crime detection, in: Pacific Visualization Symposium (PacificVis), IEEE, 2011, pp. 203–210.
780
- [8] M. L. Huang, J. Liang, Q. V. Nguyen, A visualization approach for frauds detection in financial market, in: 13th International Conference Information Visualisation, IEEE, 2009, pp. 197–202.
- 785 [9] A. Bigelow, C. Nobre, M. Meyer, A. Lex, Origraph: interactive network wrangling, in: 2019 IEEE Conference on Visual Analytics Science and Technology (VAST), IEEE, 2019, pp. 81–92.
- [10] T. Fujiwara, T. Crnovrsanin, K.-L. Ma, Concise provenance of interactive network analysis, Visual Informatics 2 (4) (2018) 213–224.
- 790 [11] S. Van den Elzen, J. J. Van Wijk, Multivariate network exploration and presentation: From detail to overview via selections and aggregations, IEEE Transactions on Visualization and Computer Graphics 20 (12) (2014) 2310–2319.

- 795 [12] W. Didimo, L. Giamminonni, G. Liotta, F. Montecchiani, D. Pagliuca, A visual analytics system to support tax evasion discovery, *Decision Support Systems* 110 (2018) 71–83.
- [13] W. Didimo, F. Giacchè, F. Montecchiani, Kojaph: visual definition and exploration of patterns in graph databases, in: *International Symposium on Graph Drawing and Network Visualization*, Springer, 2015, pp. 272–278.
- 800 [14] P.-Y. Koenig, F. Zaidi, D. Archambault, Interactive searching and visualization of patterns in attributed graphs, in: *Proceedings of Graphics Interface 2010*, Canadian Information Processing Society, 2010, pp. 113–120.
- [15] D. H. Chau, C. Faloutsos, H. Tong, J. I. Hong, B. Gallagher, T. Eliassi-Rad, Graphite: A visual query system for large graphs, in: *Data Mining Workshops, 2008. ICDMW'08. IEEE International Conference on*, IEEE, 2008, pp. 963–966.
- 805 [16] R. Pienta, A. Tamersoy, A. Endert, S. Navathe, H. Tong, D. H. Chau, Visage: Interactive visual graph querying, in: *Proceedings of the International Working Conference on Advanced Visual Interfaces*, ACM, 2016, pp. 272–279.
- 810 [17] Z. Shen, K.-L. Ma, T. Eliassi-Rad, Visual analysis of large heterogeneous social networks by semantic and structural abstraction, *IEEE Transactions on Visualization and Computer Graphics* 12 (6) (2006) 1427–1439.
- [18] E. Horvitz, Principles of mixed-initiative user interfaces, in: *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, ACM, 1999, pp. 159–166.
- 815 [19] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, M. Streit, C. Tominski, Guidance or no guidance? a decision tree can help, in: *EuroVis Workshop on Visual Analytics (EuroVA)*. The Eurographics Association, 2018.
- 820 [20] D. Ceneda, T. Gschwandtner, S. Miksch, C. Tominski, Guided visual exploration of cyclical patterns in time-series, in: *Visualization in Data Science*,

IEEE Digital Library, 2018, posterpräsentation: Visualization in Data Science (VDS at IEEE VIS 2018), Berlin; 2018-10-22.

- 825 [21] S. Miksch, W. Aigner, A matter of time: Applying a data–users–tasks design triangle to visual analytics of time-oriented data, *Computers & Graphics* 38 (2014) 286–290.
- [22] T. Munzner, A nested model for visualization design and validation, *IEEE Transactions on Visualization and Computer Graphics* 15 (6) (2009) 921–928.
- 830 [23] M. P. Timmer, E. Dietzenbacher, B. Los, R. Stehrer, G. J. de Vries, An illustrated user guide to the world input–output database: the case of global automotive production, Vol. 23, pp. 575–605.
- [24] ÖNACE, Önace 2008 - struktur, accessed: 2019-07-20.
- [25] W. Wien, Sabina - info - datenbanken, accessed: 2019-06-10.
- 835 [26] W. Aigner, S. Miksch, H. Schumann, C. Tominski, *Visualization of time-oriented data*, Springer Science & Business Media, 2011.
- [27] O. Fritz, R. Kurzmann, G. Zakarias, G. Streicher, et al., Constructing regional input-output tables for austria, *Austrian Economic Quarterly* 1 (2003) 23–39.
- 840 [28] M. Harrower, C. A. Brewer, Colorbrewer. org: an online tool for selecting colour schemes for maps, *The Cartographic Journal* 40 (1) (2003) 27–37.
- [29] M. Ghoniem, J.-D. Fekete, P. Castagliola, A comparison of the readability of graphs using node-link and matrix-based representations, in: *IEEE Symposium on Information Visualization*, IEEE, 2004, pp. 17–24.
- 845 [30] R. Keller, C. M. Eckert, P. J. Clarkson, Matrices or node-link diagrams: which visual representation is better for visualising connectivity models?, *Information Visualization* 5 (1) (2006) 62–76.

- [31] A. Cockburn, A. Karlson, B. B. Bederson, A review of overview + detail, zooming, and focus + context interfaces, *ACM Computing Surveys (CSUR)* 41 (1) (2009) 2.
- 850
- [32] W. Willett, J. Heer, M. Agrawala, Scented widgets: Improving navigation cues with embedded visualizations, *IEEE Transactions on Visualization and Computer Graphics* 13 (6) (2007) 1129–1136.
- [33] C. Ware, *Information visualization: perception for design*, Elsevier, 2012.
- 855 [34] M. Okoe, R. Jianu, S. G. Kobourov, Node-link or adjacency matrices: Old question, new insights, *IEEE Transactions on Visualization and Computer Graphics*.
- [35] D. Holten, J. J. Van Wijk, A user study on visualizing directed edges in graphs, in: *Proceedings of the SIGCHI conference on human factors in computing systems*, ACM, 2009, pp. 2299–2308.
- 860
- [36] D. Archambault, H. C. Purchase, Mental map preservation helps user orientation in dynamic graphs, in: *International Symposium on Graph Drawing*, Springer, 2012, pp. 475–486.
- [37] T. Isenberg, P. Isenberg, J. Chen, M. Sedlmair, T. Möller, A systematic review on the practice of evaluating visualization, *IEEE Transactions on Visualization and Computer Graphics* 19 (12) (2013) 2818–2827.
- 865
- [38] S. Kriglstein, M. Pohl, Choosing the Right Sample? Experiences of Selecting Participants for Visualization Evaluation, in: W. Aigner, P. Rosenthal, C. Scheidegger (Eds.), *EuroVis Workshop on Reproducibility, Verification, and Validation in Visualization (EuroRV3)*, The Eurographics Association, 2015.
- 870
- [39] N. Andrienko, G. Andrienko, *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.

- 875 [40] G. Klein, Seeing what others don't: The remarkable ways we gain insights, PublicAffairs, 2013.
- [41] Microsoft, Excel, office.microsoft.com/en-us/excel/ (accessed: 2018-10-06).
- [42] E. Wall, L. M. Blaha, L. Franklin, A. Endert, Warning, bias may occur: A proposed approach to detecting cognitive bias in interactive visual analytics, in: 2017 IEEE Conference on Visual Analytics Science and Technology (VAST), IEEE, 2017, pp. 104–115.
- 880