User Tasks & Analysis

CS 7450 - Information Visualization
September 18, 2013
John Stasko

What for?

- In order to build better visualizations, we need to understand what people might use them for
  - What tasks do they want to accomplish?

An Example

- search vs. browsing

- Value of Vis day:
  - Exploratory data analysis
  - Identifying better questions
  - Understanding, awareness, context, trust

Browsing vs. Search

- Important difference in activities
- Appears that information visualization may have more to offer to browsing

- But…browsing is a softer, fuzzier activity
- So, how do we articulate utility?
  - Maybe describe when it's useful
  - When is browsing useful?

Browsing

- Useful when
  - Good underlying structure so that items close to one another can be inferred to be similar
  - Users are unfamiliar with collection contents
  - Users have limited understanding of how system is organized and prefer less cognitively loaded method of exploration
  - Users have difficulty verbalizing underlying information need
  - Information is easier to recognize than describe

Lin '97

Thought

- Maybe infovis isn't about answering questions or solving problems… hmmm
- Maybe it's about asking better questions

Tasks

- OK, but browsing and search are very high level
- Let's be more specific…

Example from Earlier

Questions:
Which cereal has the most/least potassium?
Is there a relationship between potassium and fiber?
If so, are there any outliers?
Which manufacturer makes the healthiest cereals?

	A	B	C	D
1	Cereal	Manufacturer	Fiber	Potassium
2	100% Bran	N	10	280
3	100% Natural Bran	Q	2	135
4	All-Bran	K	9	320
5	All-Bran with Extra Fiber	K	14	330
6	Almond Delight	R	1	0
7	Apple Cinnamon Cheeric	G	1.5	70
8	Bran Chex	R	4	125
9	Bran Flakes	P	5	190
10	Cap'n'Crunch	Q	0	35
11	Cheerios	G	2	105
12	Cocoa Puffs	G	0	55
13	Corn Chex	R	0	25
14	Corn Flakes	K	1	35
15	Count Chocula	G	0	65
16	Cracklin' Oat Bran	K	4	160
17	Cream of Wheat (Quick)	N	1	0
18	Crispy Wheat & Raisins	G	2	120
19	Double Chex	R	1	80
20	Froot Loops	K	1	30
21	Frosted Flakes	K	1	25
22	Fruit & Fibre Dates, Wal	P	5	200
23	Fruitful Bran	K	5	190
24	Fruity Pebbles	P	0	25
25	Golden Grahams	G	0	45
26	Grape Nuts Flakes	P	3	85
27	Honey Nut Cheerios	G	1.5	90
28	Honey-comb	P	0	35
29	Just Right Fruit & Nut	K	2	95
30	Life	Q	2	95
31	Lucky Charms	G	0	55
32	Maypo	A	0	95
33	Muesli Raisins, Dates, &	R	3	170
34	Multi-Grain Cheerios	G	2	90
35	Nutri-Grain Almond-Rais	K	3	130
36	Nutri-grain Wheat	K	3	90
37	Oatmeal Raisin Crisp	G	1.5	120
38	Post Nat. Raisin Bran	P	6	260
39	Product 19	K	1	45
40	Quaker Oatmeal	Q	2.7	110
41	Raisin Bran	K	5	240
42	Raisin Nut Bran	G	2.5	140
43	Rice Krispies	K	0	35
44	Shredded Wheat	N	3	95
45	Shredded Wheat 'n'Bran	N	4	140
46	Shredded Wheat spoon	N	3	120
47	Smacks	K	1	40
48	Special K	K	1	55
49	Strawberry Fruit Wheats	N	3	90
50	Total Corn Flakes	G	0	35
51	Total Raisin Bran	G	4	230
52	Total Whole Grain	G	3	110
53	Trix	G	0	25
54	Wheaties	G	3	110
55	Wheaties Honey Gold	G	1	60

Exercise

- What are the (types of) tasks being done here?
- Can you think of others?
  - Let's develop a list

Task Taxonomies

- Number of different ones exist, important to understand what process they focus on
  - Creating an artifact
  - Human tasks
  - Tasks using visualization system
  - …

User Tasks

- Wehrend & Lewis created a low-level, domain independent taxonomy of user tasks in visualization environments
- Eleven basic actions
  - identify, locate, distinguish, categorize, cluster, distribution, rank, compare within relations, compare between relations, associate, correlate

Wehrend & Lewis
Vis '90

Another Perspective

- Shneiderman proposed task × data type taxonomy to understand what people do with visualization
- Mantra: "Overview first, zoom and filter, then details on demand"
  - Design paradigm for infovis systems

Shneiderman
VL '96

Taxonomy

- Data Types
  1. 1D
  2. 2D
  3. 3D
  4. Temporal
  5. ND
  6. Tree
  7. Network
- Tasks
  1. Overview
  2. Zoom
  3. Filter
  4. Details-on-demand
  5. Relate
  6. History
  7. Extract

Another Task Taxonomy

- Amar, Eagan, & Stasko – InfoVis ’05

Background

- Use "commercial tools" class assignment from this class
- Students generate questions to be answered using commercial infovis systems
- Data sets:

Domain	Data cases	Attributes	Questions Generated
Cereals	78	15	107
Mutual funds	987	14	41
Cars	407	10	153
Films	1742	10	169
Grocery surveys	5164	8	126

- Generated 596 total analysis tasks

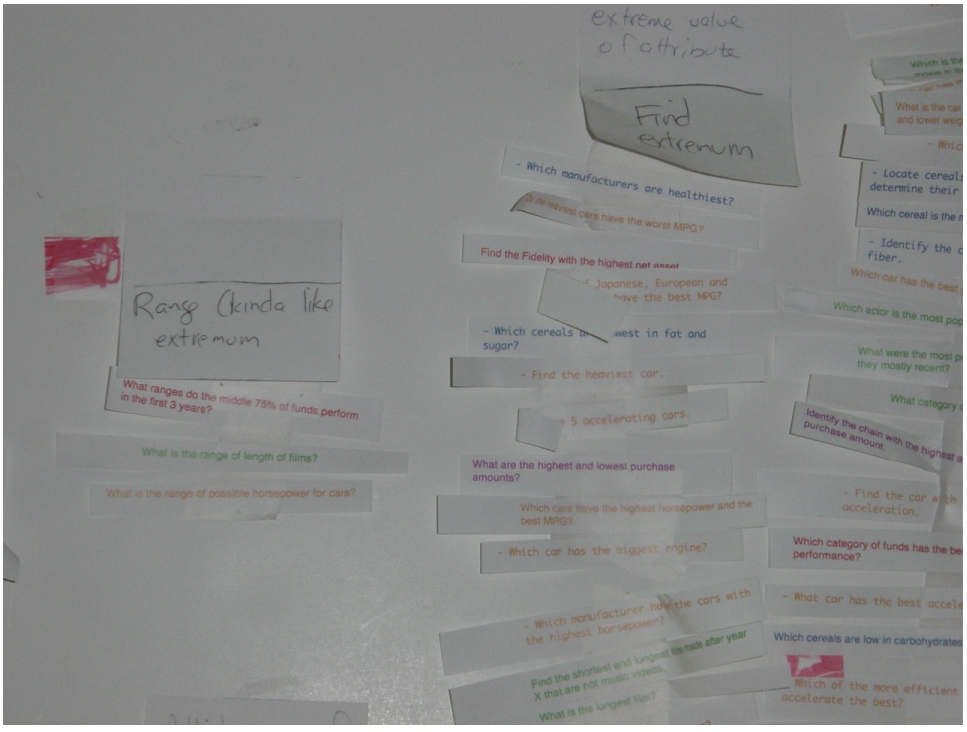
extreme value of attribute
Find extremum
Range (kinda like extremum
- Which manufacturers are healthiest?
Find the Fidelity with the highest net asset
- Japanese, European and
- Which cereals
sugar?
- Find the heaviest car.
What ranges do the middle 75% of funds perform in the first 3 years?
What is the range of length of films?
What is the range of possible horsepower for cars?
What are the highest and lowest purchase amounts?
Which cars have the highest horsepower and the best MPG?
- Which car has the biggest engine?
Which manufacturer has the cars with the highest horsepower?
Find the shortest and longest film made after year
What is the longest film?
Which cereals are low in carbohydrates
Which of the more efficient accelerate the best?
- Find the car with acceleration.
Which category of funds has the be performance?
- What car has the best accele
Identify the chain with the highest purchase amount.
What were the most p they mostly recent?
What category


SAVE
Please do not erase
Organizational operations (filter, cluster, categ.)
vs.
transformation ops (avg, count, ...)
1-1 n-1
find extremum
evidence for extremum
aggregate
Filter
Cluster
Categorize
look-up value(s)
(find entity)
Simple Comparison
Distribution
Clustering
Sort
Range
Correlation
Browse
Outliers
Categorical Correlation
Extremum of X
count
Filter
What is the car w/ highest MPG?
What director has won the most awards?
What film has Robin Williams' film has release date ↑?
-> case(s) at extreme of range of attr(s)
-> What is mpg of Audi TT?
-> How long is Gone w/ Wind?
-> Given case(s), what are some attributes?
-> compare scalar values
-> compare the MPG of American and Japanese cars.
-> Characterization of distribution of an attribute over set of cases
-> What is the distribution of carbs in cereals?
-> What is age distribution of shoppers?
-> Find clusters of attribute values in a set of cases
-> Are there groups of cereals w/ similar fat/cal/sugar?
-> Is there a cluster of typical film length?
-> Range of an attribute in a set of cases
-> What is the range of the length of films?
-> What is the range of horsepower of cars?
-> Order set of cases by some metric
-> Order the cars by weight
-> Rank the cereals by calories
How many films are 127 min long?
How many manufacturers of cars are there?
-> Find cardinality of set of cases
-> What kellogs cereals have >3g fiber?
-> What comedies have won awards?
-> What funds under performed S&P 500?
read vs. correlate

Terminology

- *Data case* – An entity in the data set
- *Attribute* – A value measured for all data cases
- *Aggregation function* – A function that creates a numeric representation for a set of data cases (eg, average, count, sum)

1. Retrieve Value

General Description:
Given a set of specific cases, find attributes of those cases.

Examples:
- What is the mileage per gallon of the Audi TT?
- How long is the movie Gone with the Wind?

2. Filter

General Description:
Given some concrete conditions on attribute values, find data cases satisfying those conditions.

Examples:
- What Kellogg's cereals have high fiber?
- What comedies have won awards?
- Which funds underperformed the SP-500?

3. Compute Derived Value

General Description:
Given a set of data cases, compute an aggregate numeric representation of those data cases.

Examples:
- What is the gross income of all stores combined?
- How many manufacturers of cars are there?
- What is the average calorie content of Post cereals?

4. Find Extremum

General Description:

Find data cases possessing an extreme value of an attribute over its range within the data set.

Examples:

- What is the car with the highest MPG?
- What director/film has won the most awards?
- What Robin Williams film has the most recent release date?

5. Sort

General Description:

Given a set of data cases, rank them according to some ordinal metric.

Examples:

- Order the cars by weight.
- Rank the cereals by calories.

6. Determine Range

General Description:
Given a set of data cases and an attribute of interest, find the span of values within the set.

Examples:
- What is the range of film lengths?
- What is the range of car horsepowers?
- What actresses are in the data set?

7. Characterize Distribution

General Description:
Given a set of data cases and a quantitative attribute of interest, characterize the distribution of that attribute's values over the set.

Examples:
- What is the distribution of carbohydrates in cereals?
- What is the age distribution of shoppers?

8. Find Anomalies

General Description:
Identify any anomalies within a given set of data cases with respect to a given relationship or expectation, e.g. statistical outliers.

Examples:
- Are there any outliers in protein?
- Are there exceptions to the relationship between horsepower and acceleration?

9. Cluster

General Description:
Given a set of data cases, find clusters of similar attribute values.

Examples:
- Are there groups of cereals w/ similar fat/calories/sugar?
- Is there a cluster of typical film lengths?

10. Correlate

General Description:
Given a set of data cases and two attributes, determine useful relationships between the values of those attributes.

Examples:
- Is there a correlation between carbohydrates and fat?
- Is there a correlation between country of origin and MPG?
- Do different genders have a preferred payment method?
- Is there a trend of increasing film length over the years?

Discussion/Reflection

- Compound tasks
  - "Sort the cereal manufacturers by average fat content"
    Compute derived value; Sort
  - "Which actors have co-starred with Julia Roberts?"
    Filter; Retrieve value

Discussion/Reflection

- What questions were left out?
  - Basic math
    "Which cereal has more sugar, Cheerios or Special K?"
    "Compare the average MPG of American and Japanese cars."
  - Uncertain criteria
    "Does cereal (X, Y, Z…) sound tasty?"
    "What are the characteristics of the most valued customers?"
  - Higher-level tasks
    "How do mutual funds get rated?"
    "Are there car aspects that Toyota has concentrated on?"
  - More qualitative comparison
    "How does the Toyota RAV4 compare to the Honda CRV?"
    "What other cereals are most similar to Trix?"

Concerns/Limitations

- InfoVis tools may have influenced students' questions
- Graduate students as group being studied
  - How about professional analysts?
- Subjective – Not an exact science
- Data was really quantitative so may get a different set of tasks for relational/graph data
  - See Lee et al, BELIV '06

Contributions

- Set of grounded low-level analysis tasks
- Potential use of tasks as a language/vocabulary for comparing and evaluating infovis systems

Another Perspective

- Taxonomy proposed
- "…used specifically for multidimensional visualizations, taking into account the generic objectives that a user has when using such techniques to perform exploratory analyses as a previous step of statistical analysis."

Valiati et al
BELIV '06

Task Taxonomy

- 7 tasks in 2 categories
  - User goals
    - Identify – Find, discover new information
    - Determine – Calculate, define a precise value
    - Compare – Compare data & values
    - Infer – Infer knowledge, generate hypotheses
    - Locate – Search and identify information
  - Intermediate level tasks to support analysis
    - Visualize – Represent the data a certain way
    - Configure – Normalize, filter, reorder, etc.

Valiati et al BELIV '06

More Details

- Each task has "parameters"
  - Identify
    - clusters
    - correlations
    - categories
    - properties
    - patterns
    - characteristics
    - thresholds
    - similarities
    - differences
    - dependencies
    - uncertainties
    - variations

Interaction

- User goals and tasks carried out through interaction with visualization
  - The interactive dialog helps people explore

Interaction Framework

- Organized along *user intent*
- 7 categories
  - Select
  - Explore
  - Reconfigure
  - Encode
  - Abstract/elaborate
  - Filter
  - Connect

More to come later on interaction day

Yi et al *TVCG* '07

Interactive Dynamics

- "taxonomy of interactive dynamics that contribute to successful analytic dialogues"
  - part interaction, part task

Data and View Specification	**Visualize** data by choosing visual encodings.
	Filter out data to focus on relevant items.
	Sort items to expose patterns.
	Derive values or models from source data.
View Manipulation	**Select** items to highlight, filter, or manipulate them.
	Navigate to examine high-level patterns and low-level detail.
	Coordinate views for linked, multidimensional exploration.
	Organize multiple windows and workspaces.
Process and Provenance	**Record** analysis histories for revisitation, review, and sharing.
	Annotate patterns to document findings.
	Share views and annotations to enable collaboration.
	Guide users through analysis tasks or stories.

Heer & Shneiderman *CACM* '12

Can InfoVis Be More?

- Is InfoVis helping people enough?
- What do we need to do to provide even more value?

Providing Better Analysis

- Combine computational analysis approaches such as data mining with infovis
  - Too often viewed as competitors in past
- Each has something to contribute

Shneiderman
Information Visualization `02

Issues

- Issues influencing the design of discovery tools:
  - Statistical Algorithms vs. Visual data presentation
  - Hypothesis testing vs. exploratory data analysis
- Pro's and Con's?

Differing Views

- Hypothesis testing
  - Advocates:
    By stating hypotheses up front, limit variables and sharpens thinking, more precise measurement
  - Critics:
    Too far from reality, initial hypotheses bias toward finding evidence to support it
- Exploratory Data Analysis
  - Advocates:
    Find the interesting things this way, we now have computational capabilities to do them
  - Skeptics:
    Not generalizable, everything is a special case, detecting statistical replationships does not infer cause and effect

Recommendations

- Integrate data mining and information visualization
- Allow users to specify what they are seeking
- Recognize that users are situated in a social context
- Respect human responsibility

Another Question?

- Are the visualizations helping with exploratory analysis enough?
- Are they attempting to accomplish the right goals?

Status Quo Limitations

- Current Information Visualization systems inadequately support decision making:
  - Limited Affordances
  - Predetermined Representations
  - Decline of Determinism in Decision-Making
- "Representational primacy" versus "Analytic primacy"

Amar & Stasko
TVCG '05

Goal: High-Level Tasks

- Complex decision-making, especially under uncertainty
- Learning a domain
- Identifying the nature of trends
- Predicting the future
- …

Analytic Gaps

- Analytic gaps – "obstacles faced by visualizations in facilitating higher-level analytic tasks, such as decision making and learning."
  - Worldview Gap
  - Rationale Gap

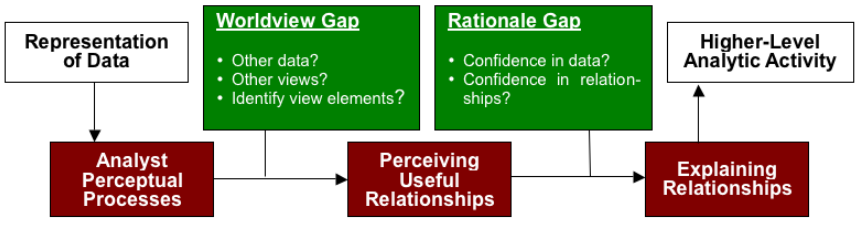

Knowledge Precepts

- For narrowing these gaps
  - Worldview-Based Precepts ("Did we show the right thing to the user?")
    - Determine Domain Parameters
    - Expose Multivariate Explanation
    - Facilitate Hypothesis Testing
  - Rationale-Based Precepts ("Will the user believe what they see?")
    - Expose Uncertainty
    - Concretize Relationships
    - Expose Cause and Effect

Application of Precepts

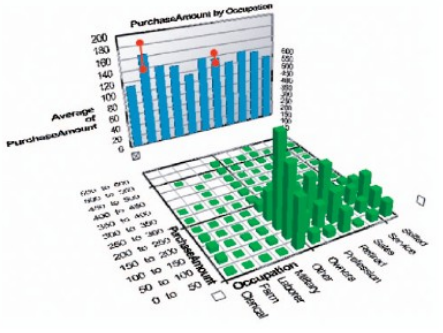

Fig. 2. Error bars (which we have added in red) would be a simple way to increase confidence in the degree of difference between two aggregations. (Picture taken from the SeeIt system by Visible Decisions, Inc.)

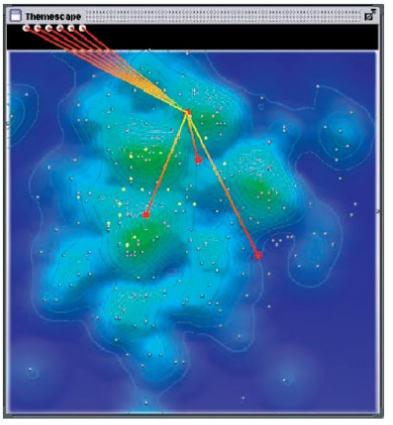

Fig. 3. This themescape variation allows documents with missing metadata, shown as dots in the upper black region, to participate in analysis, such as the reference relationship shown. (Picture courtesy of Nicholas Diakopoulos.)

Application of Precepts

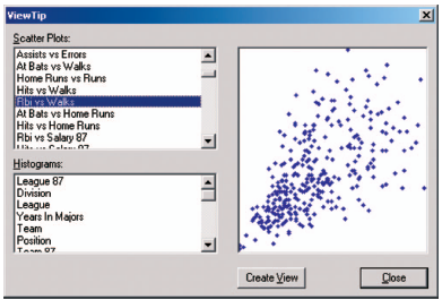

Fig. 4. The View Tips in SpotFire Pro 4.0 allow users to quickly examine possible sources of correlation for further examination.

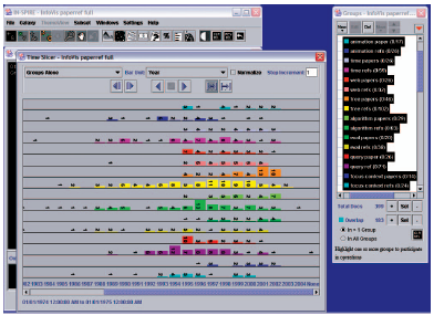

Fig. 5. IN-SPIRE uses horizontal scrolling to navigate time slices of user-defined content groups. (Picture produced at and provided with permission of Pacific Northwest National Laboratory, which is managed and operated by the Battelle Memorial Institute on behalf of the US Department of Energy.)

Put Them Together

- Combine the ideas:
  - Use computational, statistical analysis more
  - Cater to the user's analytic reasoning needs
- And put together with infovis

- Leads to…

Visual Analytics

- "The science of analytical reasoning facilitated by interactive visual interfaces"
- Combines
  - Data analysis
  - Infovis
  - Analytical reasoning
- Grew from view that infovis was neglecting these other aspects
  - True?

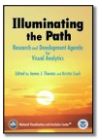

Thomas & Cook
Illuminating the Path

Visual Analytics

- Grew from stimulus in the homeland security area
  - Need for better data analysis methods
  - Really big data
- Topic for entire week later in term…

Related Detour

- Your projects are "design studies"
  - Problem-driven visualization research
  - Assist clients with data who want to understand it better
  - Design and build visualization system

- How do you do it well?

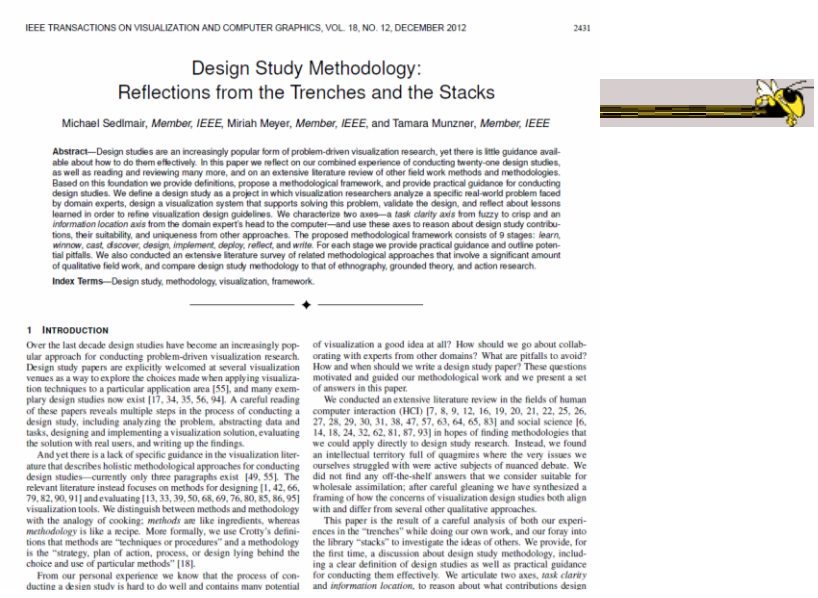

IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS, VOL. 18, NO. 12, DECEMBER 2012 2431

Design Study Methodology:
Reflections from the Trenches and the Stacks

Michael Sedlmair, *Member, IEEE*, Miriah Meyer, *Member, IEEE*, and Tamara Munzner, *Member, IEEE*

Abstract—Design studies are an increasingly popular form of problem-driven visualization research, yet there is little guidance available about how to do them effectively. In this paper we reflect on our combined experience of conducting twenty-one design studies, as well as reading and reviewing many more, and on an extensive literature review of other field work methods and methodologies. Based on this foundation we provide definitions, propose a methodological framework, and provide practical guidance for conducting design studies. We 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. We characterize two axes—a *task clarity axis* from fuzzy to crisp and an *information location axis* from the domain expert's head to the computer—and use these axes to reason about design study contributions, their suitability, and uniqueness from other approaches. The proposed methodological framework consists of 9 stages: *learn, winnow, cast, discover, design, implement, deploy, reflect*, and *write*. For each stage we provide practical guidance and outline potential pitfalls. We also conducted an extensive literature survey of related methodological approaches that involve a significant amount of qualitative field work, and compare design study methodology to that of ethnography, grounded theory, and action research.

Index Terms—Design study, methodology, visualization, framework.

1 Introduction

Over the last decade design studies have become an increasingly popular approach for conducting problem-driven visualization research. Design study papers are explicitly welcomed at several visualization venues as a way to explore the choices made when applying visualization techniques to a particular application area [55], and many exemplary design studies now exist [17, 34, 35, 56, 94]. A careful reading of these papers reveals multiple steps in the process of conducting a design study, including analyzing the problem, abstracting data and tasks, designing and implementing a visualization solution, evaluating the solution with real users, and writing up the findings.

And yet there is a lack of specific guidance in the visualization literature that describes holistic methodological approaches for conducting design studies—currently only three paragraphs exist [49, 55]. The relevant literature instead focuses on methods for designing [1, 42, 66, 79, 82, 90, 91] and evaluating [13, 33, 39, 50, 68, 69, 76, 80, 85, 86, 95] visualization tools. We distinguish between methods and methodology with the analogy of cooking; *methods* are like ingredients, whereas *methodology* is like a recipe. More formally, we use Crotty's definitions that methods are "techniques or procedures" and a methodology is the "strategy, plan of action, process, or design lying behind the choice and use of particular methods" [18].

From our personal experience we know that the process of conducting a design study is hard to do well and contains many potential of visualization a good idea at all? How should we go about collaborating with experts from other domains? What are pitfalls to avoid? How and when should we write a design study paper? These questions motivated and guided our methodological work and we present a set of answers in this paper.

We conducted an extensive literature review in the fields of human computer interaction (HCI) [7, 8, 9, 12, 16, 19, 20, 21, 22, 25, 26, 27, 28, 29, 30, 31, 38, 47, 57, 63, 64, 65, 83] and social science [6, 14, 18, 24, 32, 62, 81, 87, 93] in hopes of finding methodologies that we could apply directly to design study research. Instead, we found an intellectual territory full of quagmires where the very issues we ourselves struggled with were active subjects of nuanced debate. We did not find any off-the-shelf answers that we consider suitable for wholesale assimilation; after careful gleaning we have synthesized a framing of how the concerns of visualization design studies both align with and differ from several other qualitative approaches.

This paper is the result of a careful analysis of both our experiences in the "trenches" while doing our own work, and our foray into the library "stacks" to investigate the ideas of others. We provide, for the first time, a discussion about design study methodology, including a clear definition of design studies as well as practical guidance for conducting them effectively. We articulate two axes, *task clarity* and *information location*, to reason about what contributions design

Reflects on 21 design studies from 3 authors & reviewing others

Definition

- "A design study is 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."

Problem Suitability

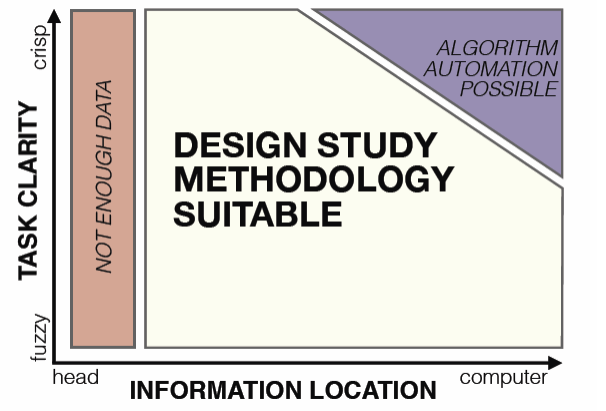

Framework

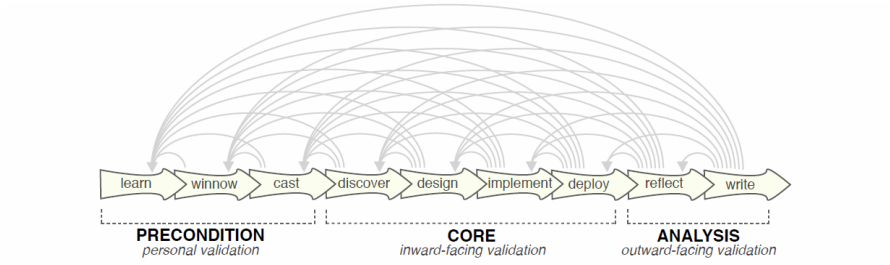

Fig. 2. Nine-stage design study methodology framework classified into three top-level categories. While outlined as a linear process, the overlapping stages and gray arrows imply the iterative dynamics of this process.

Considerations

- Practical
  - Data: Does data exist, is it enough, can you get it?
  - Engagement: How much time do they and you have for the project? How much time can you spend in their environment?
- Intellectual
  - Problem: Is there a vis research question lurking?
  - Need: Is there a real need or are existing approaches good enough?
  - Task: Are you addressing a real task? How long will need persist? How many people care?
- Interpersonal
  - What is your rapport with clients?

Pitfalls

- 32 pitfalls to design study projects listed, organized by framework phase
  - Examples
    - No real data available
    - No need for vis, problem can be automated
    - Nonrapid prototyping
    - Premature and insufficient deployment

Project Proposals & HW 2

- Back on Monday

HW 3

- HW description has more about D3
- New alternate assignment
  - Design, not programming

Upcoming

- InfoVis Systems & Toolkits
  - Reading:
    Viegas et al '07 (ManyEyes)
    Bostock et al '11 (D3)
- Commercial Systems & Tools
  - Reading:
    Spenke & Beilken '00

References

- Spence & CMS texts
- All referred to papers