

Perspicanads

An annotated portfolio of how visual tools can extend the human mind

W. Bradford Paley, Digital Image Design Incorporated

June, 2021

Per-spi-ca-cious: having a ready insight into and understanding of subtle or complex things

Per-spic-u-ous: (of a representation) clearly expressed, easily understood; lucid

Per-spi-cand: (our back-formation) the easily-understood representation

A New Kind of Tool *Illustrates the task itself, not the tool's rules*

Knowledge work is taking a significant step forward because of a new kind of tool. Risk can be reduced; speed and quality of insights and actions can be improved; creativity, morale, and customer satisfaction can be raised; training and support costs can be almost eliminated.

These benefits result from combining hundreds of small techniques that arise from a single fundamental shift in perspective: make the tool into the task—*fit it to the cognitive steps the human mind does to accomplish a task*. We do this instead of forcing people to learn the rules of every idiosyncratic tool, then map all their scattered findings back to what they're trying to do. Every task we've studied has a spatial grounding. Even abstract thought is grounded in spatial thinking. Visuals don't just *represent* information—when they illustrate it well they *become* the information; everything we need.

“I bring evidence to show that spatial thinking, acting in the world with the things in the world, is the foundation of thought; not the entire edifice, but the foundation.” Barbara Tversky, *Mind in Motion*

“If I pursue a beam of light with the velocity c (velocity of light in a vacuum), I should observe such a beam of light as an electromagnetic field at rest though spatially oscillating.” Albert Einstein's visual thought experiment, from *Autobiographical Notes*

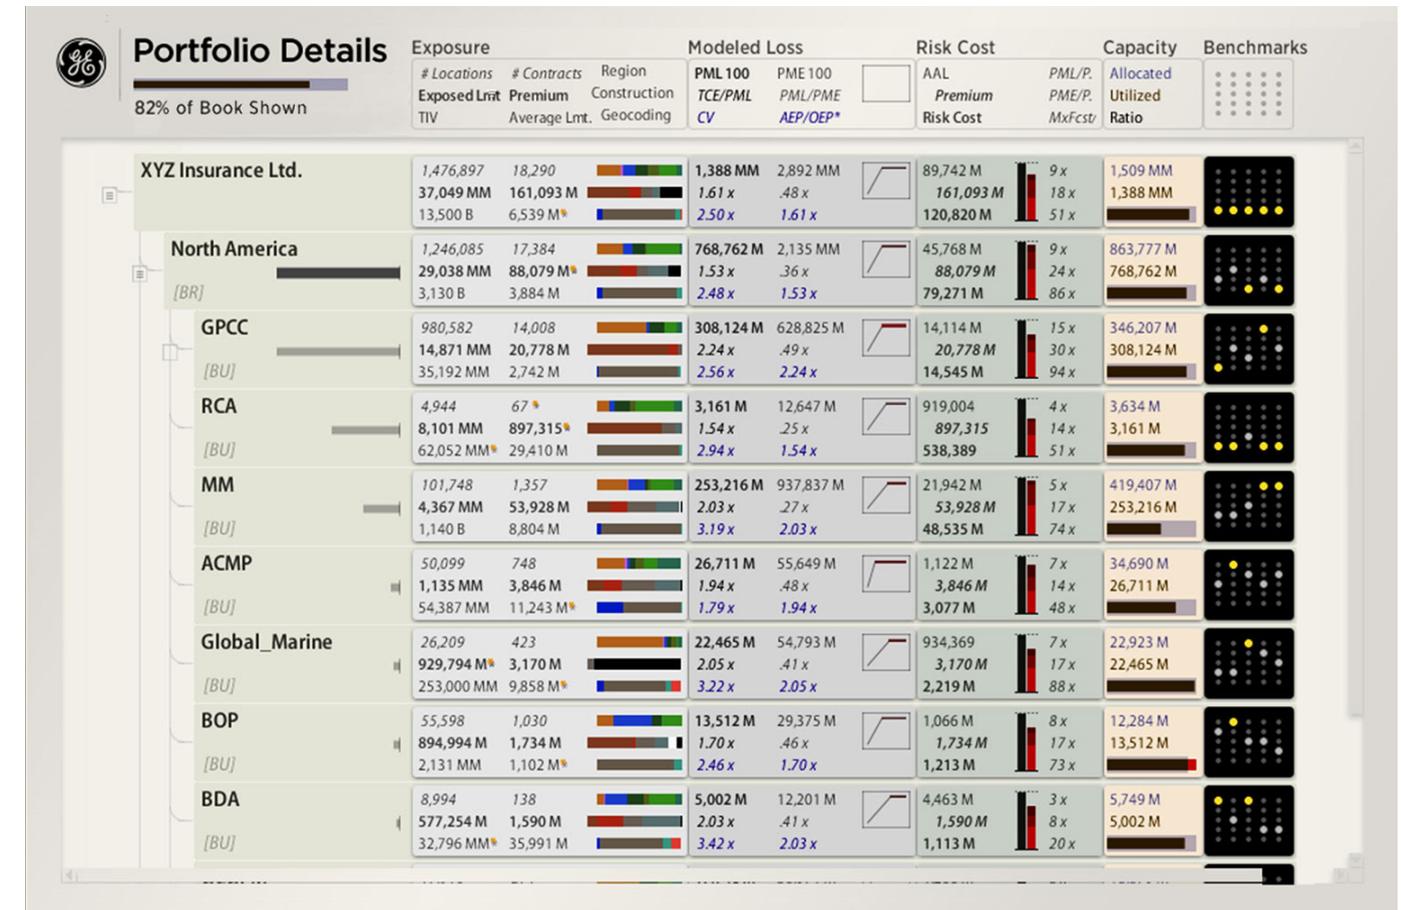
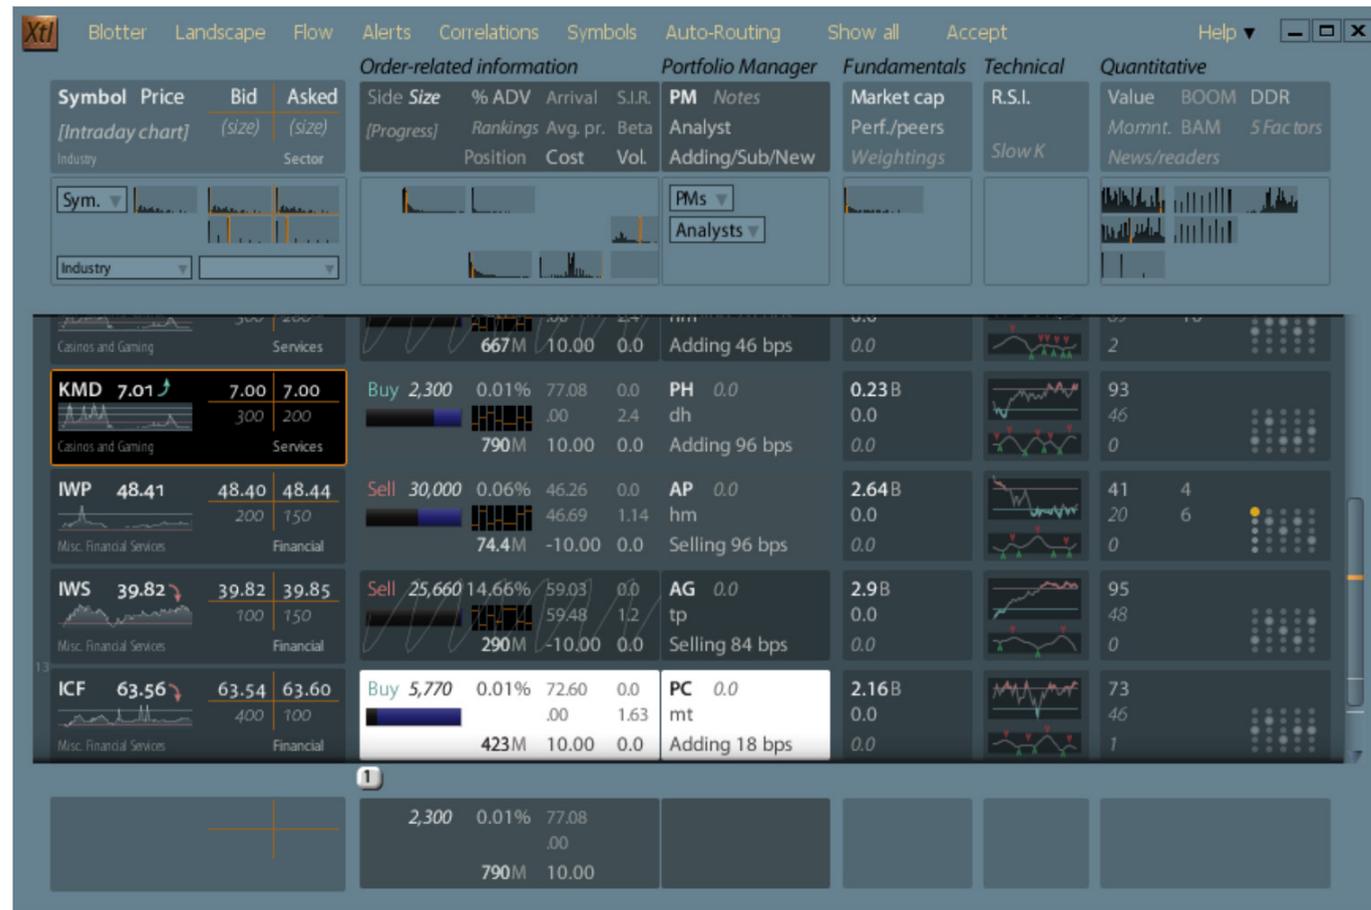
“We knew that DNA was important...But we didn't realise I think just how important it would be. Put in other words, we didn't realise that the shape would give us a clue to the replication mechanism.” Francis Crick, *BBC TV interview*

My collaborators and I have quietly deployed dozens of these tools over the last two decades. You're likely to have seen one or two at work in the New York Stock Exchange or other major financial institutions, featured in the New York Times, Nature, Geo, ID Magazine, SlashDot; or on the wall at MoMA or the Whitney Museum of American Art. It's now time to broaden their use; time to help more people.

We gather many here, in categories identifying the kinds of thought processes they support. Innovated, honed, tested, and proven; we're now stretching this approach to toolmaking by addressing the most critical human/data challenges. Chances are they're not as complex as you thought.

Dashboards *Overviews let us get our bearings and prioritize tasks*

But must be tuned: they give us the lay of the land best when fitted to specific tasks and cognitive sequences.



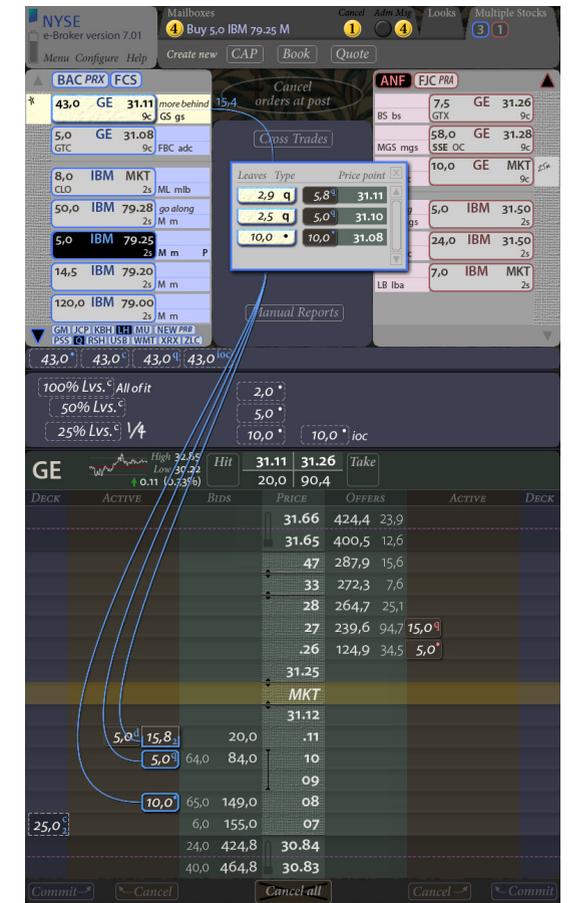
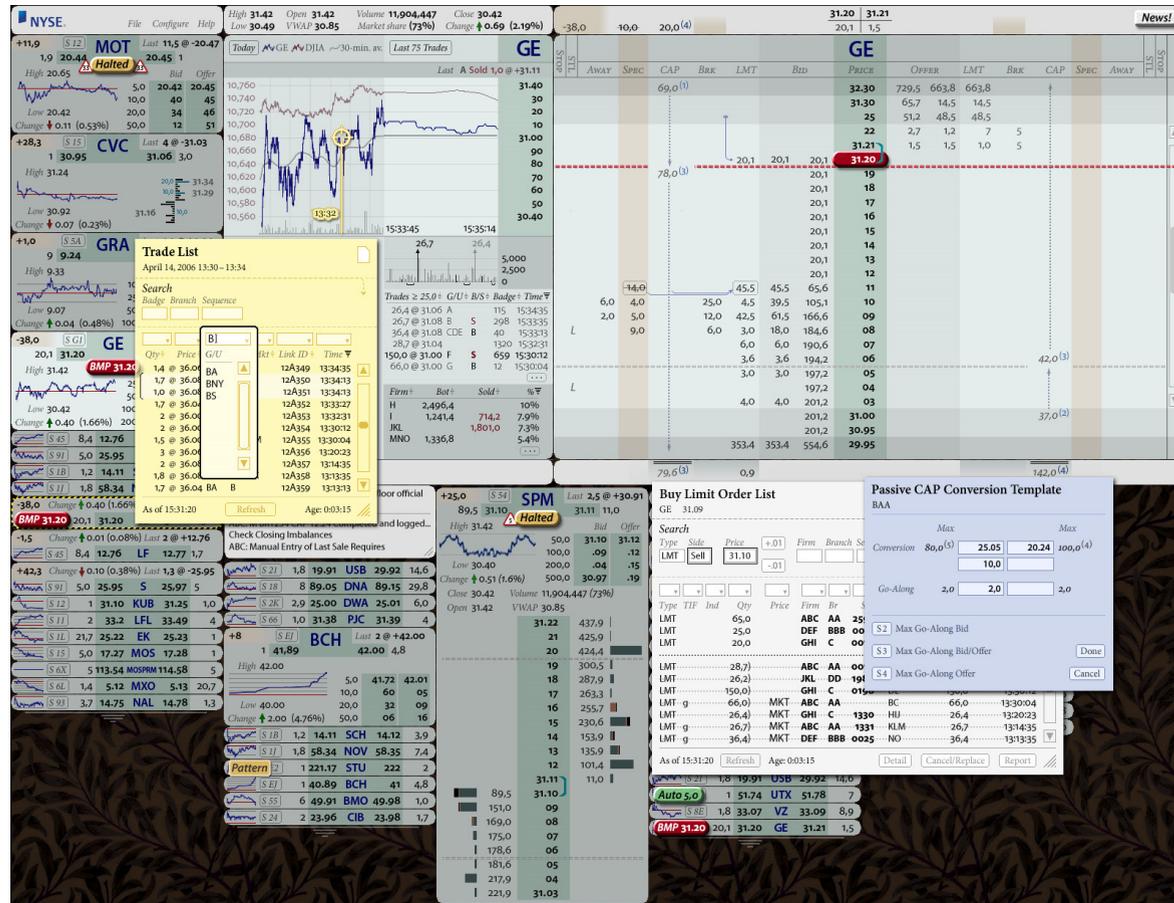
An equities trading blotter and an insurance risk-management portfolio viewer group information by intellectual context. In each, 50+ columns worth of data were arranged so no horizontal scrolling was necessary. This was absolutely critical to traders who needed to get an order working in the market within 30 seconds; before the off-screen columns were functionally completely unavailable.

In general: Spreadsheet-like grids and blotters become important decision-support tools with complete data: necessary to choose the next task—but they're rarely sufficient to complete it in an optimal manner.

Following slides: beyond dashboards and orienting oneself, visuals can support many distinct categories of thinking, but to do so their designer must know the precise operations of the business—and the mind.

Transaction Support *After the overview: focus on a single task*

We are more effective, more comfortable, and get into intellectual flow more easily when a tool exactly fits a task.



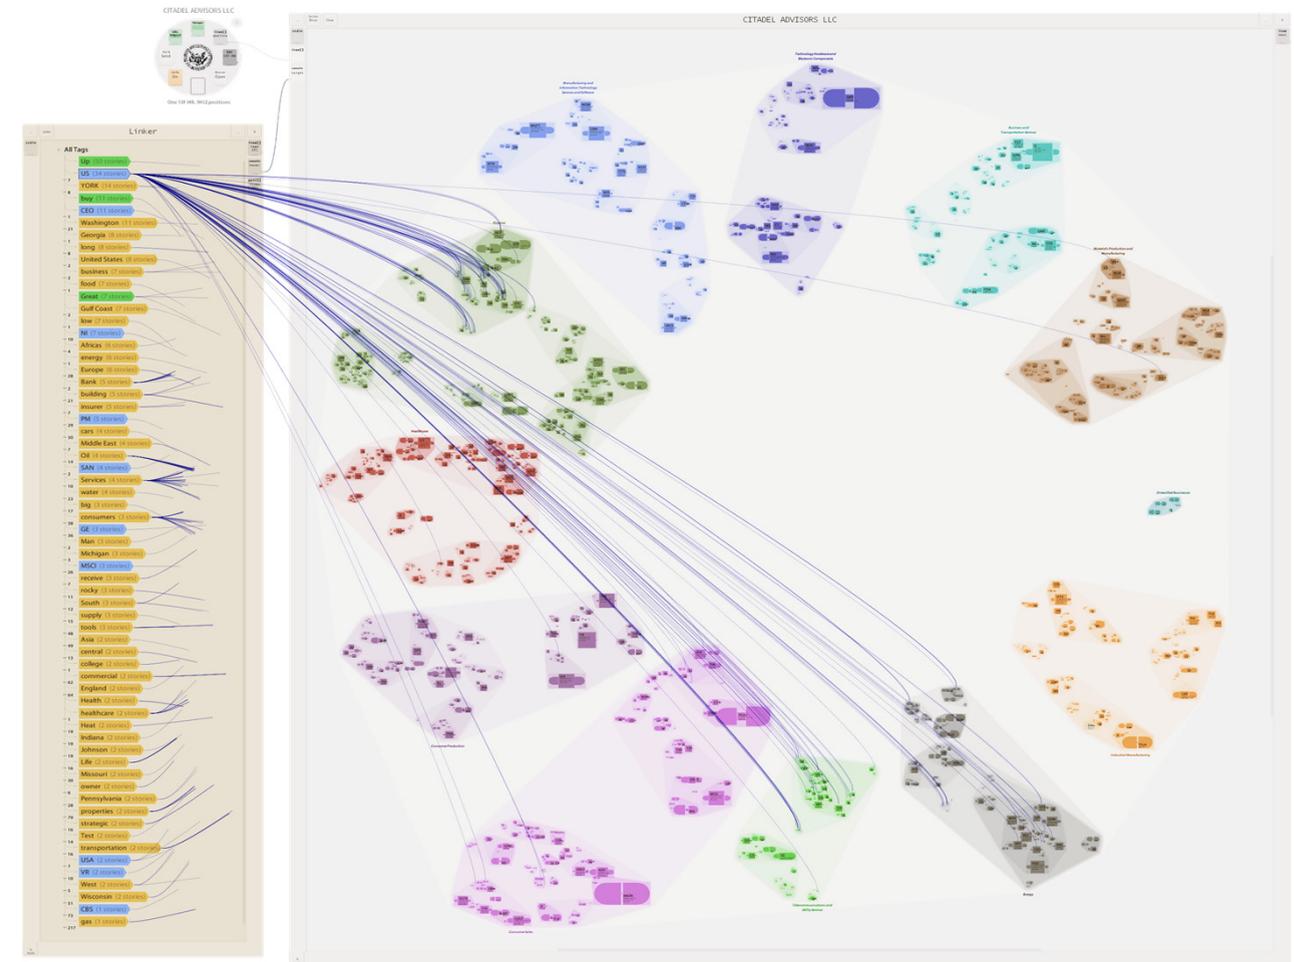
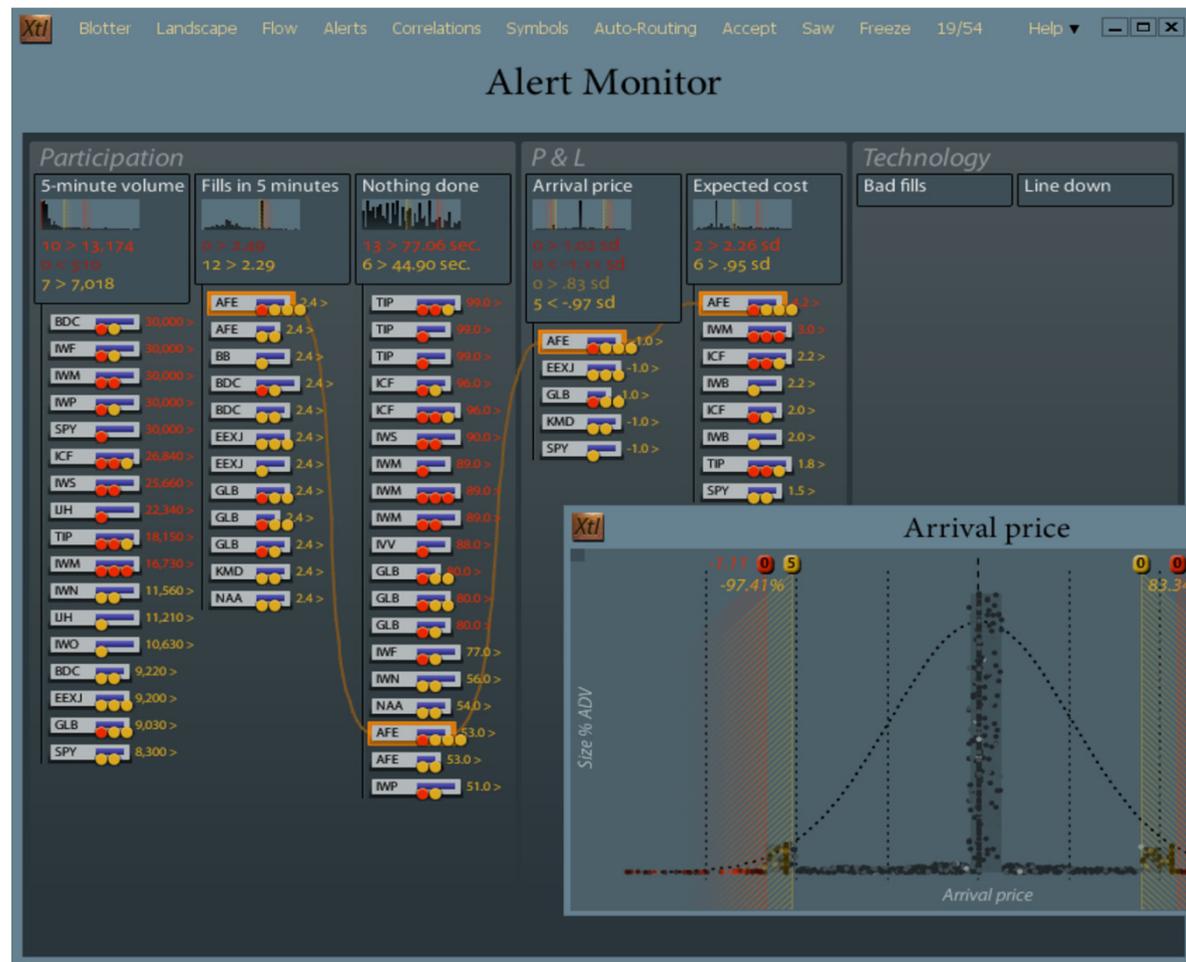
Forcing experts to do every task with a scattered set of apps means they have to integrate all the data they need—every day, every time they do the task. Our tools integrate it all a single time, during development, with microservices and interop tech—much more important: by fitting tools to a proper cognitive deconstruction of what's needed. This is done before any code is written, so the effort's properly factored out in a project's non-recurring engineering phase.

The effort of integration never even gets to, e.g., the trading desk. Left: a NYSE Specialist's workstation shows 5x as many stocks as before, and data for some 15x as many actionable insights (a better metric for display effectiveness).

Center: our Active Order Landscape ranks orders by expected cost (vertical) and footprint (horizontal), so traders work the upper-right quadrant. Better than any ordered list: they can make intelligent trade-offs. Clicking brings up a placement tool that ranks all possible venues by all relevant criteria. Right: our NYSE handheld speed broker "refills" from 10-15 seconds down to less than one: they just click a source order and drag a blueprint template into the market.

Alerting Problems aren't one-dimensional—and they can live anywhere

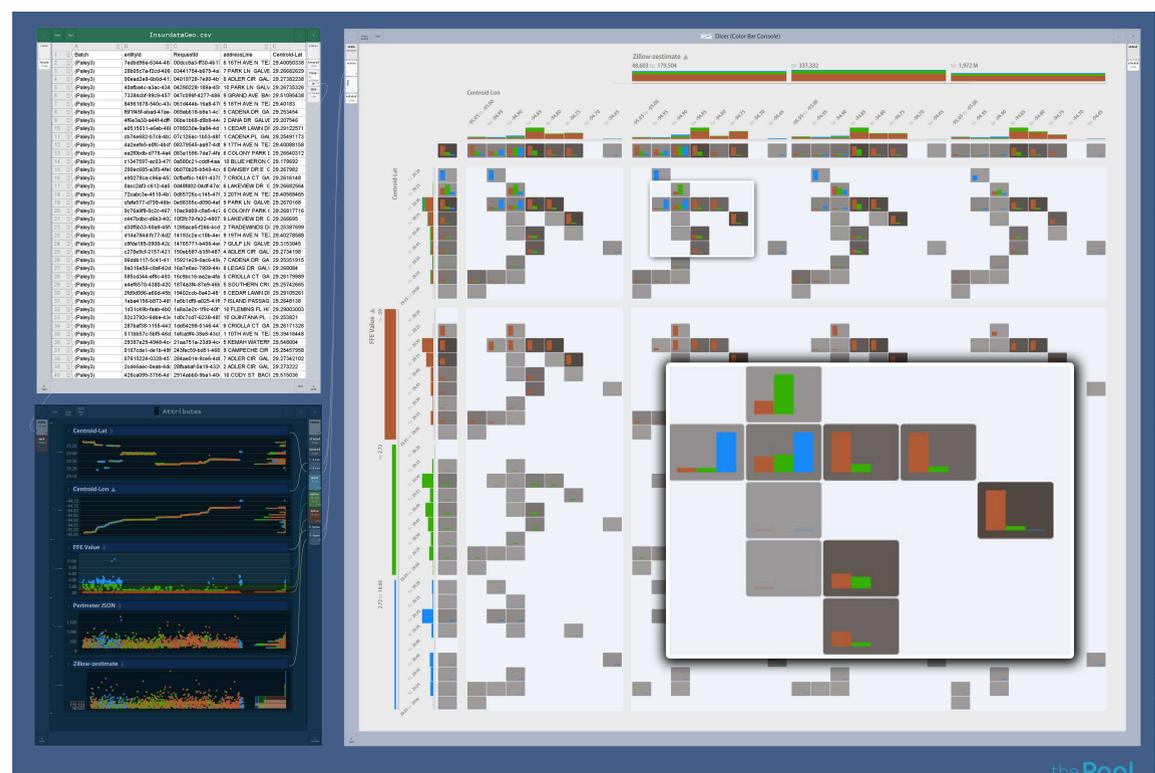
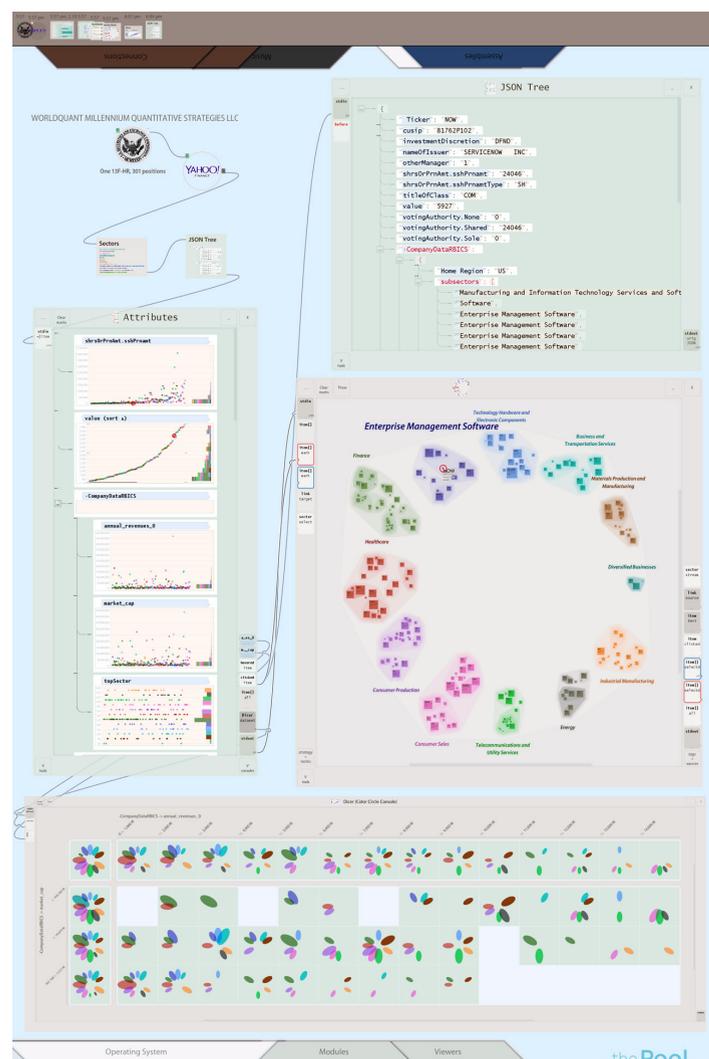
A common strategy tries to order all alerts in one list. The natural world does it differently—and far better.



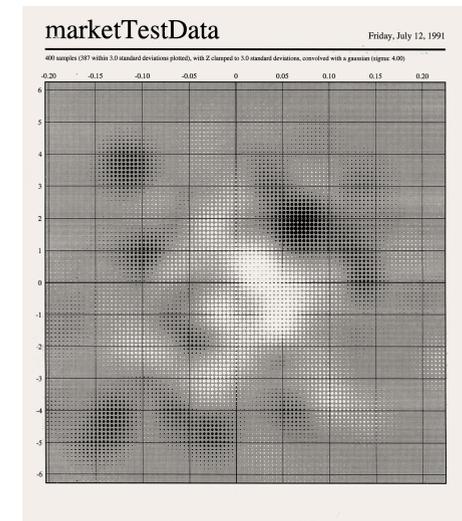
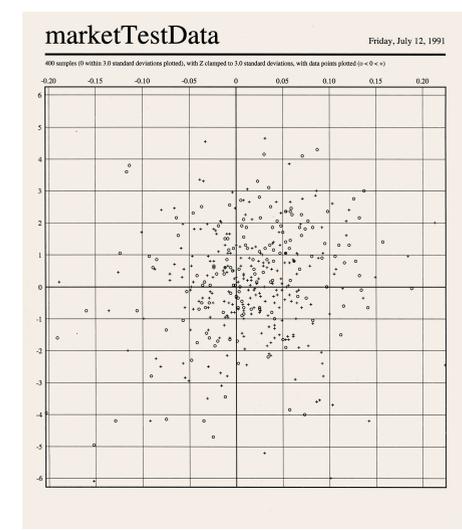
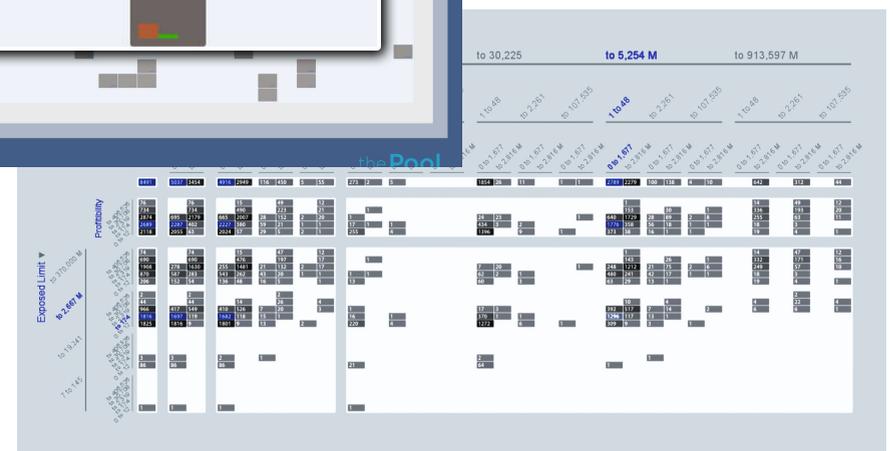
In a jungle “alerts” can be anywhere: anything unusual—but we react to them just fine. Left: some alerting needs many lists: one for each critical decision dimension. (Also: real-world data is virtually never distributed in a normal curve. Here we show real data (histogram) so we can set amber/red alert thresholds just before/after the well-populated one-dollar points. We’ve backed out the normal curve & standard deviations to show how badly they fit—use SDs and you’re just not fitting reality.) Right: a set of tags points into the known, everyday “jungle” of the expert (here, a visualization of portfolio positions sorted into hierarchical sectors). When hovered, a tag completes its connections. But even the stubs of the lines from the other tags show useful information about how applicable each is to the portfolio, and an approximation of where.

Visual Analytics *The eye supports the mind; lets it think more deeply*

Manipulating data, e.g., in multiple dimensions—and seeing it that way—lets people find higher-level patterns.



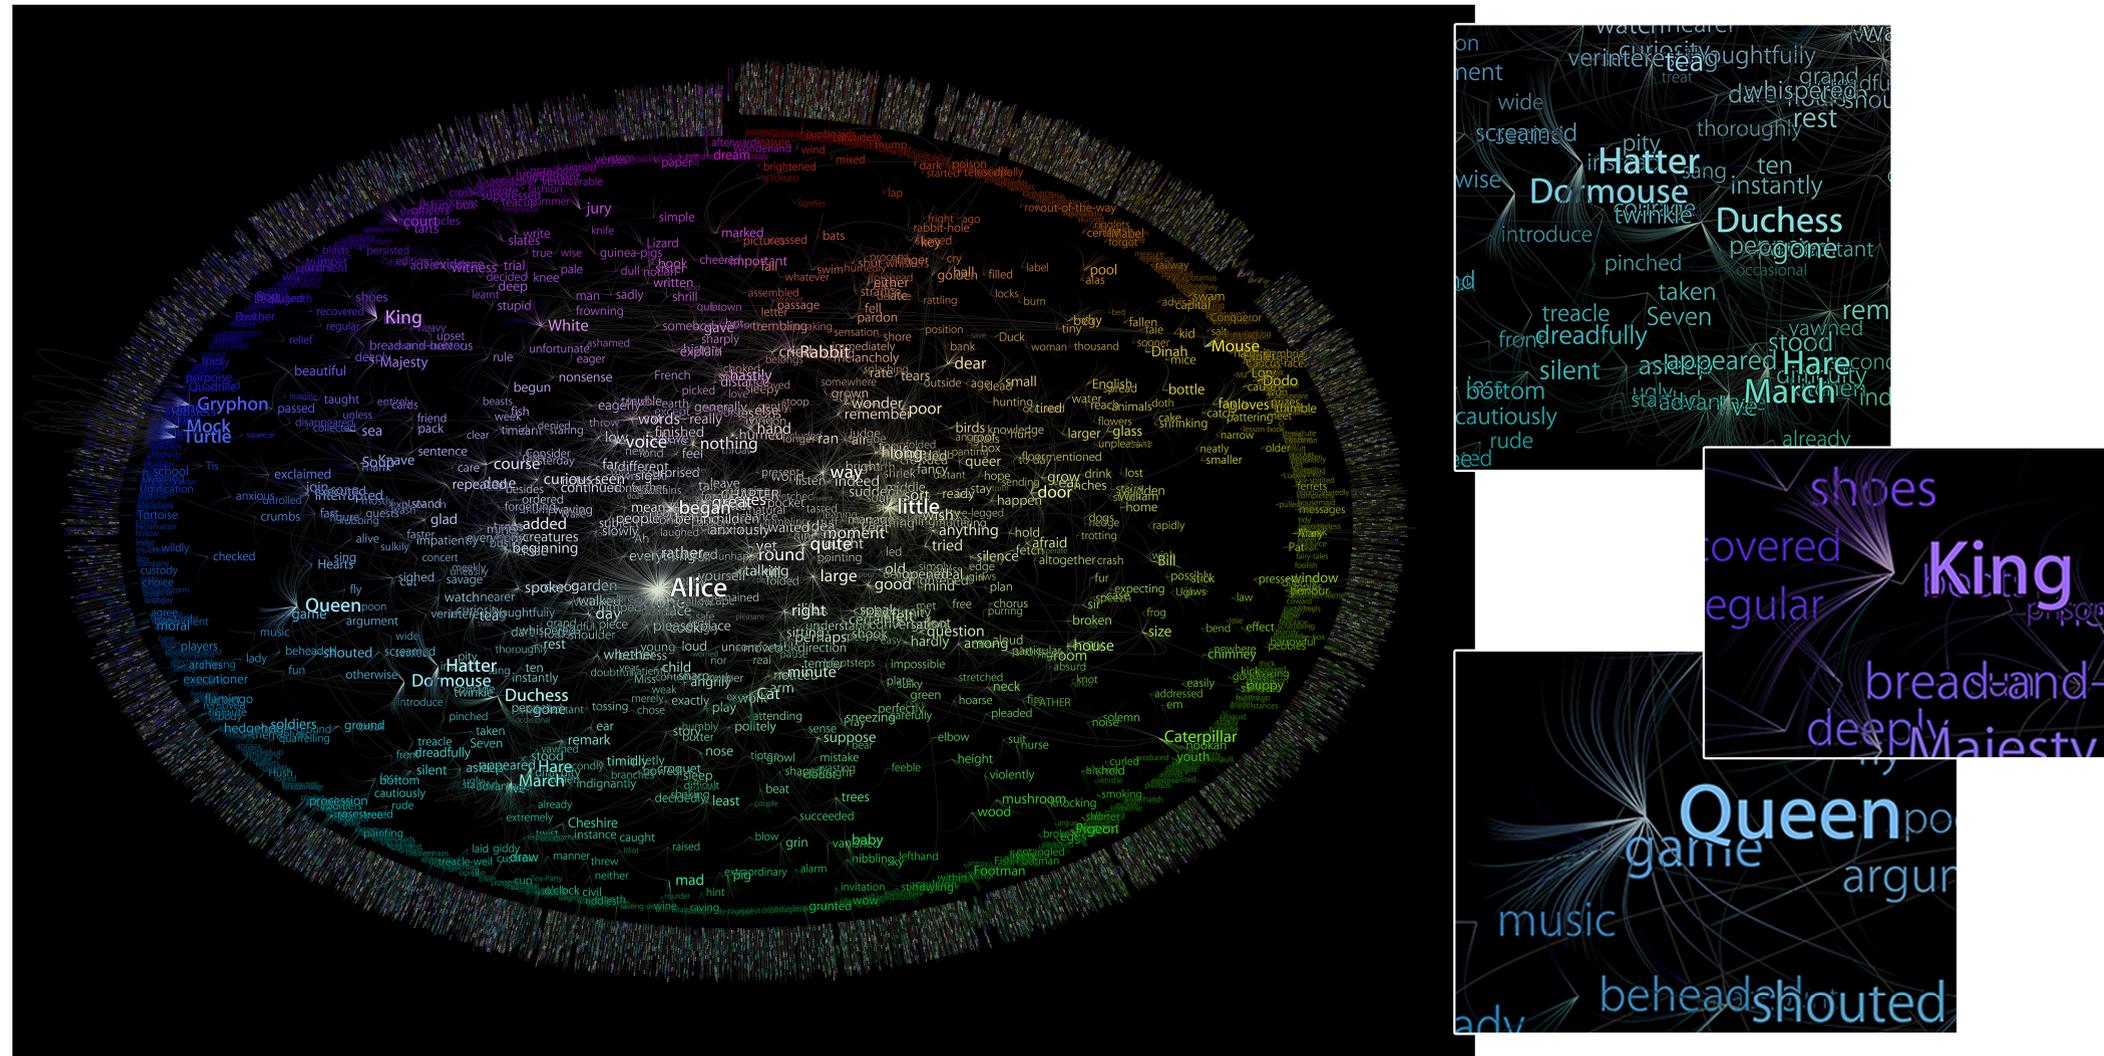
Our integration/analytcs/collaboration/presentation environment, *the Pool*, lets data and ideas flow: from input to transforms to filters to visuals and back out again. Every step in a complex process can not only be watched, but understood by others.



People can only rehearse 7 ± 2 chunks of information in short-term memory. And there's a huge cost to swapping in a new chunk; if you have to read or navigate to it you may lose two or three. But when visuals act as a kind of second-level mental cache this swap is instant. Cognitive load is reduced, leaving more of your mind free to generate insights. Left: a 13F-HR flows through our AttributeView to feed a 2D table in our Dicer, an n-D data-cube viewer. Center: a 4D table, with a bar chart in every cell—ideally it's an entire task-specific mini-report. Under: a 5D breakdown of an insurance portfolio. Right, a bit of history: we pioneered financial data visualization at Lehman Brothers from 1989 on. Before (top): try to find + marks among the circles; after: trade where this neural net's outputs convolve+sum to be white. (This image taught us to trade against the net where it's black!) Note that we use small circles on gray to create a reading that's both qualitative and quantitative: one can match the size of the circle to a key much more accurately than judging a shade of gray—especially since the eye reads brightness and color by local context, assuring only a local maximum.

Hypothesis Formation *Brainstorm with visuals—but don't commit*

The eye is fooled more easily than statistics. It's powerful—but the result of 510,000,000 years of evolutionary kludges.



Here TextArc shows distribution glyphs for each of the 2500+ individual word stems in Alice's Adventures in Wonderland. Their star-like rays point to the word's usage in the text, arranged clockwise around the arc.

This lets Structuralists form hypotheses like "Why are Hatter, Dormouse and March Hare all distributed like the word 'tea'?" This can then be verified with research (which, for textual analysis, can be as quick as reading the text right in the tool).

More interesting: why do King and Queen appear in almost all the same places— except for four thin lines emanating lower right from Queen? Textual research here show us that's what foreshadowing looks like. She doesn't actually appear yet, she's just mentioned in advance so she has a more dramatic entrance when she does.

This is not to say we shouldn't use visuals! They provide an unparalleled scaffolding for higher-level thought. But those thoughts need to be researched and tested using less idiosyncratic methods. Many of the eye's biases can be corrected for (e.g., by TextArc's multi-attribute layering, or our bright gaps in the upcoming What-if slide), but this requires a deep understanding of both eye & mind I wasn't taught as a graphic artist or coder. Critically: if we want to be rigorous, to support anything like the scientific method visuals must be tightly integrated with verification tools, as well as tools that let people take action on what they see. A stand-alone visualization is a dead end, not part of a process. One needs to be able to grab and use what one sees. Note: TextArc's layout, glyph, and layering system can be applied to any dataset where the goal is understanding thousands of individual elements (think of tags or attributes) with respect to a single linear ranking (not necessarily the tiny text here).

Causality *What will follow, as the practice of analytics evolves*

Causal models: stronger than correlation, but we're just starting to understand them. Intervening helps.

The Relative Importance of Heredity and Environment in Determining the Piebald Pattern of Guinea-Pigs

W. B. Peley, After Sewall Wright: FIG. 5, Proceedings of the National Academy of Sciences, June 1920, p. 328

The Relative Importance of Heredity and Environment in Determining the Piebald Pattern of Guinea-Pigs

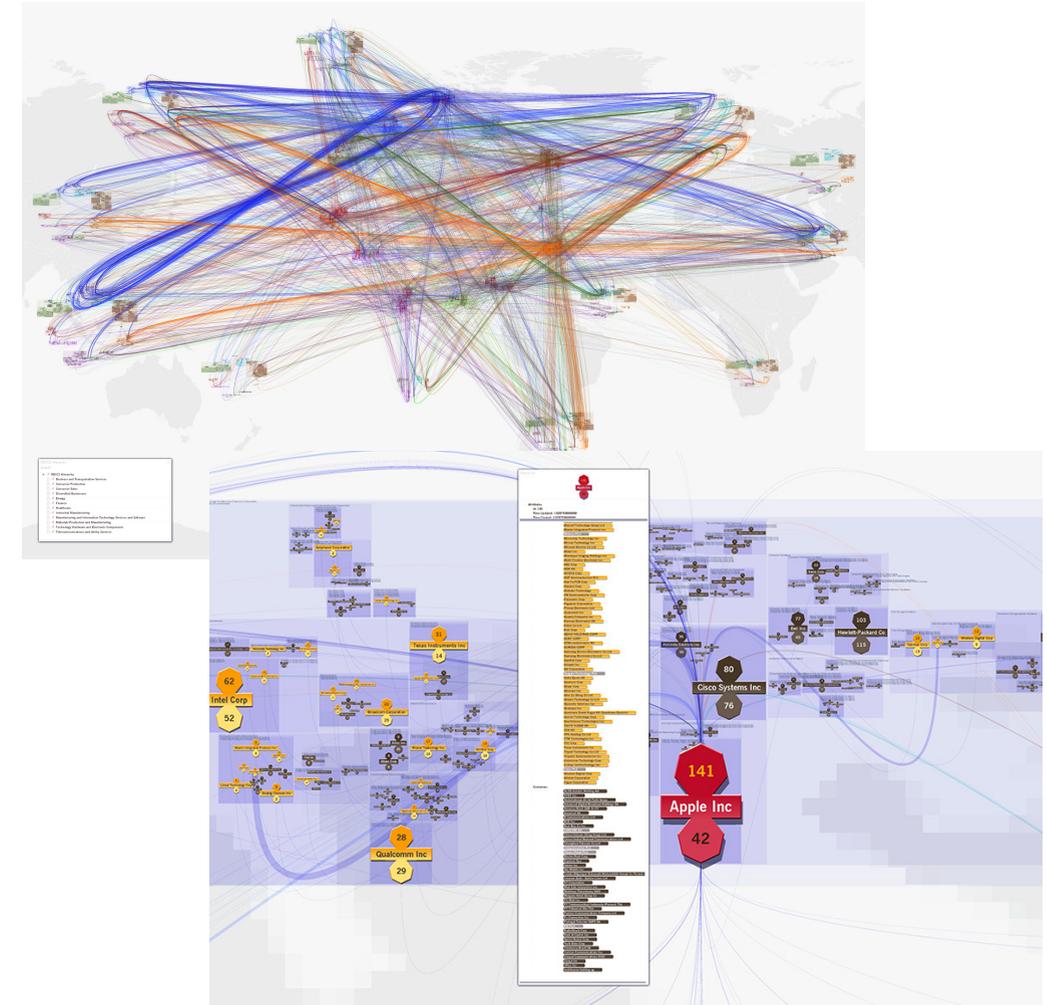
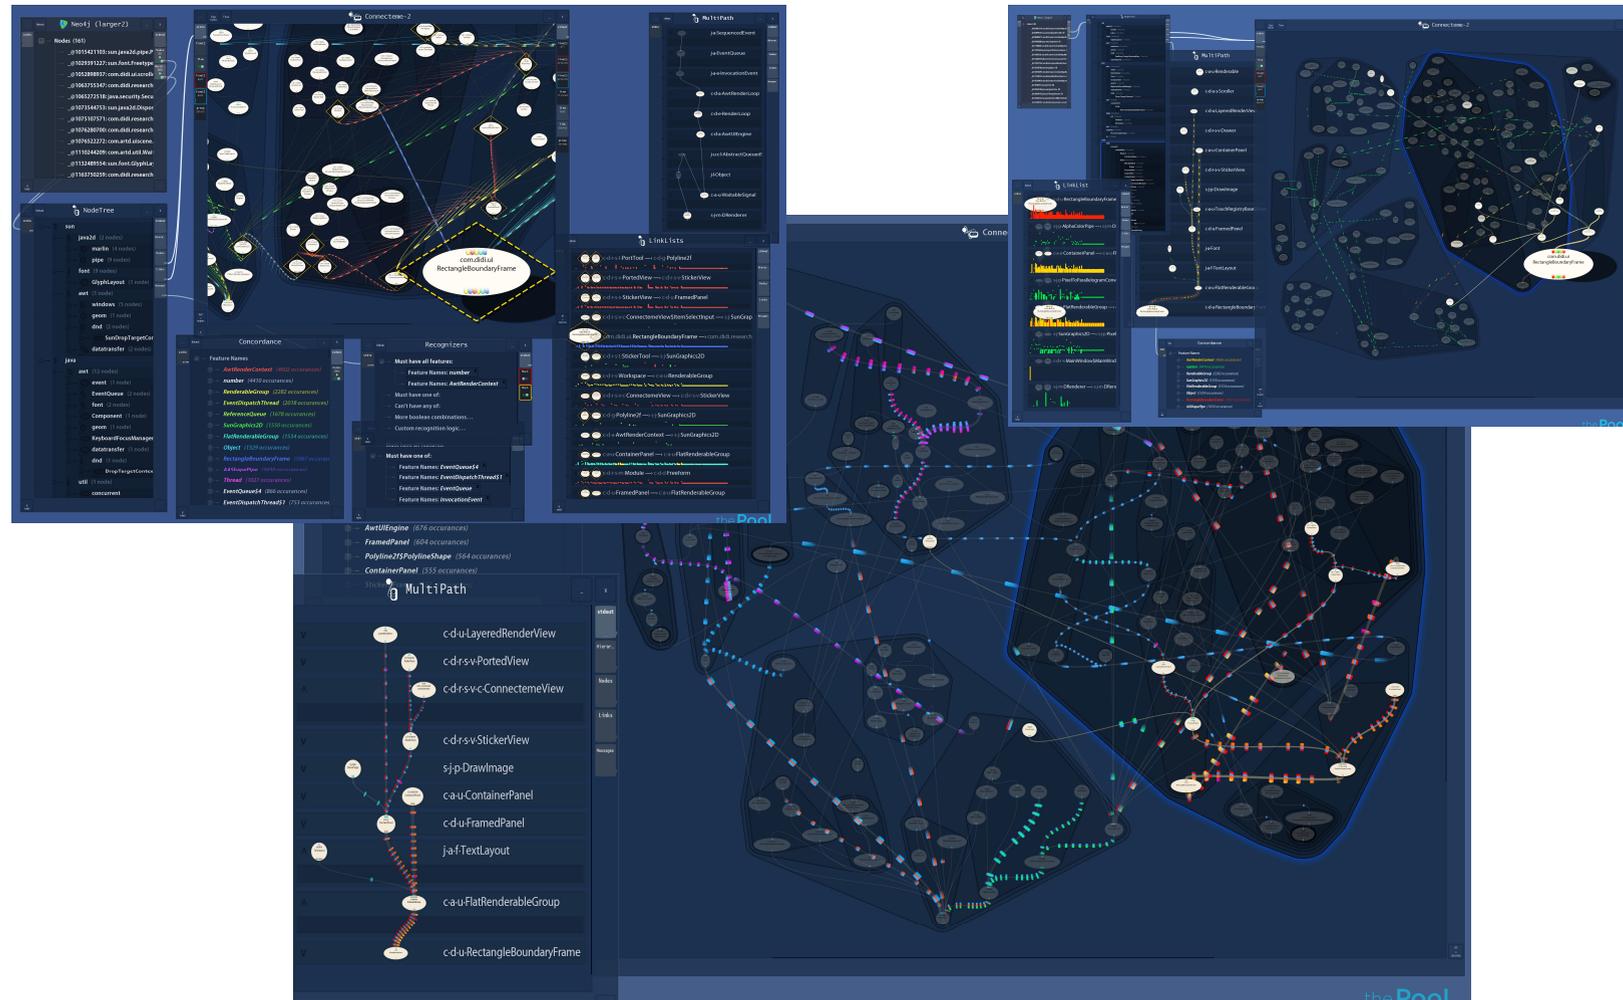
W. B. Peley, After Sewall Wright: FIG. 5, Proceedings of the National Academy of Sciences, June 1920, p. 328

Pulling "a strand of causal pearls" out of a network

Judea Pearl developed Bayesian networks, but when he found he couldn't get them to incorporate the concept of cause and effect he left them to his graduate students. Bayesian methods changed the world nonetheless, but Pearl went on to develop Causal Models—and they're beginning to change things again. Why? Killing the rooster will **not** stop the sun from rising, but most statistical models can't tell you that.

Network Understanding *Where GraphVis & metrics end up, if humans lead them*

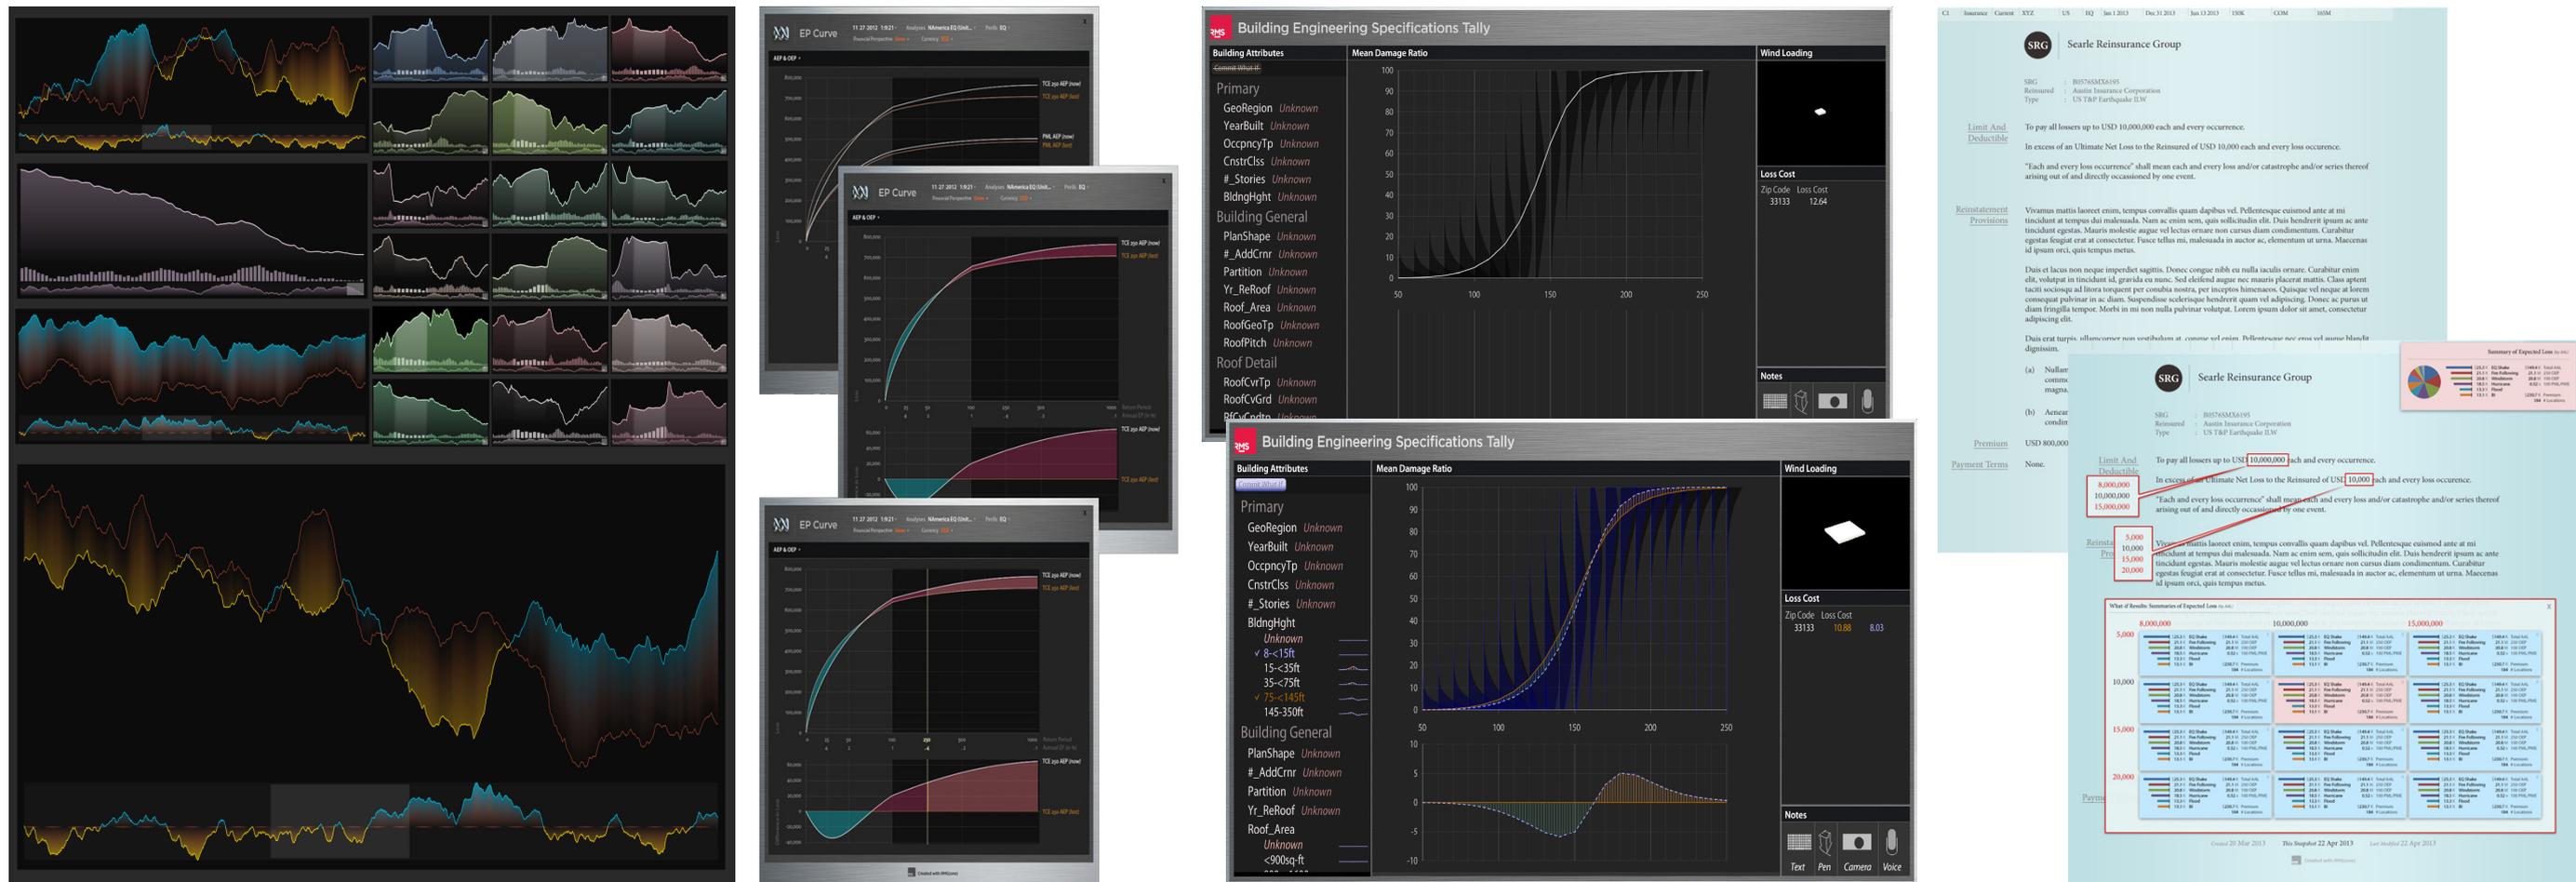
Graphs encode *every part* of a complex issue as nodes/edges. But extracting meaning is hard; a human thing.



People learn better when they can play with and organize things themselves, compared to passive viewing. This is well known in pedagogy and cognitive science. Our tools support sense-making and evoke actionable insights because people can reach in to add the meaning they know, e.g., C-suite on top, infrastructure bottom. It's hard to try to learn how a robot arranged content based on topology alone. It's an attractive fallback when semantics aren't available, and creates lovely organic-looking slides, but that's probably the rules of the layout algorithm more than the data. Domain experts put in the data's meaning, joining as intellectual cyborgs when their top-down organizing skills are supported by bottom-up algorithmic bookkeeping. To operate at the level of meaning we find experts and fit to their analytic best practices—the tools are informed by real-world processes and goals. Right: global supply chains: 6k nodes & 30k links, but still readable.

Comparisons and What-Ifs *A number isn't enough on its own*

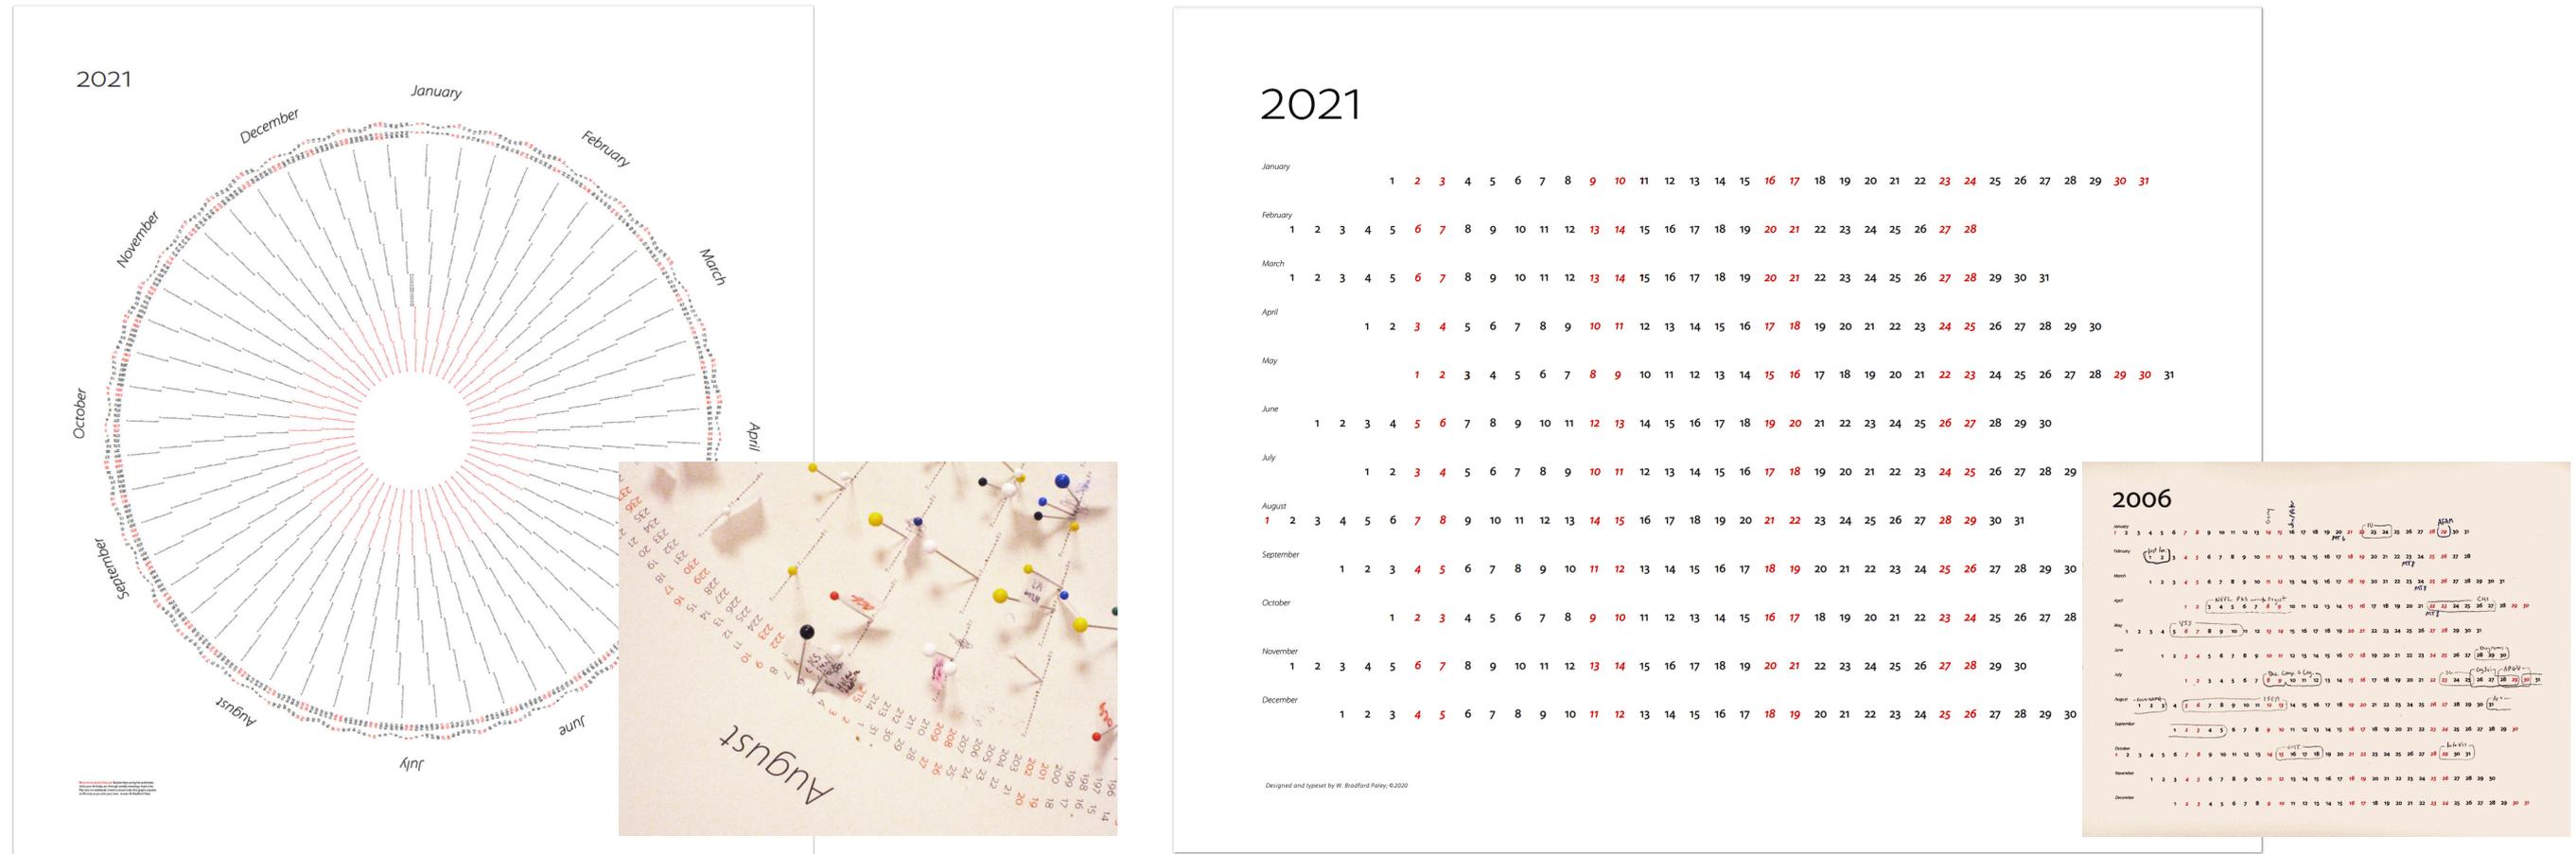
“Statistics is the art of *with respect to what.*” Good visuals make the context relevant, clear, and deep.



I worked briefly/deeply with Edward Tufte, and this simple phrase was a powerful gift. (Though we agreed to disagree on rules like his data/ink ratio; ink (his inverse proxy for simplicity) isn't the scarce resource, attention is. We try to minimize "data/think.") Left: Human eyes are biased to measure gaps perpendicular to lines; it was evolutionarily important to judge a branch's strength, an enemy's arm. But it's a serious error when reading charts. We make larger gaps brighter. Center left: It's hard to compare close lines—until we subtract them & show the difference. Center right: Why decide based on just averages when you can see full distributions everywhere along the curve? Right: contract what-ifs: right in context.

Everyday Usage *Quiet back there, please, I'm thinking*

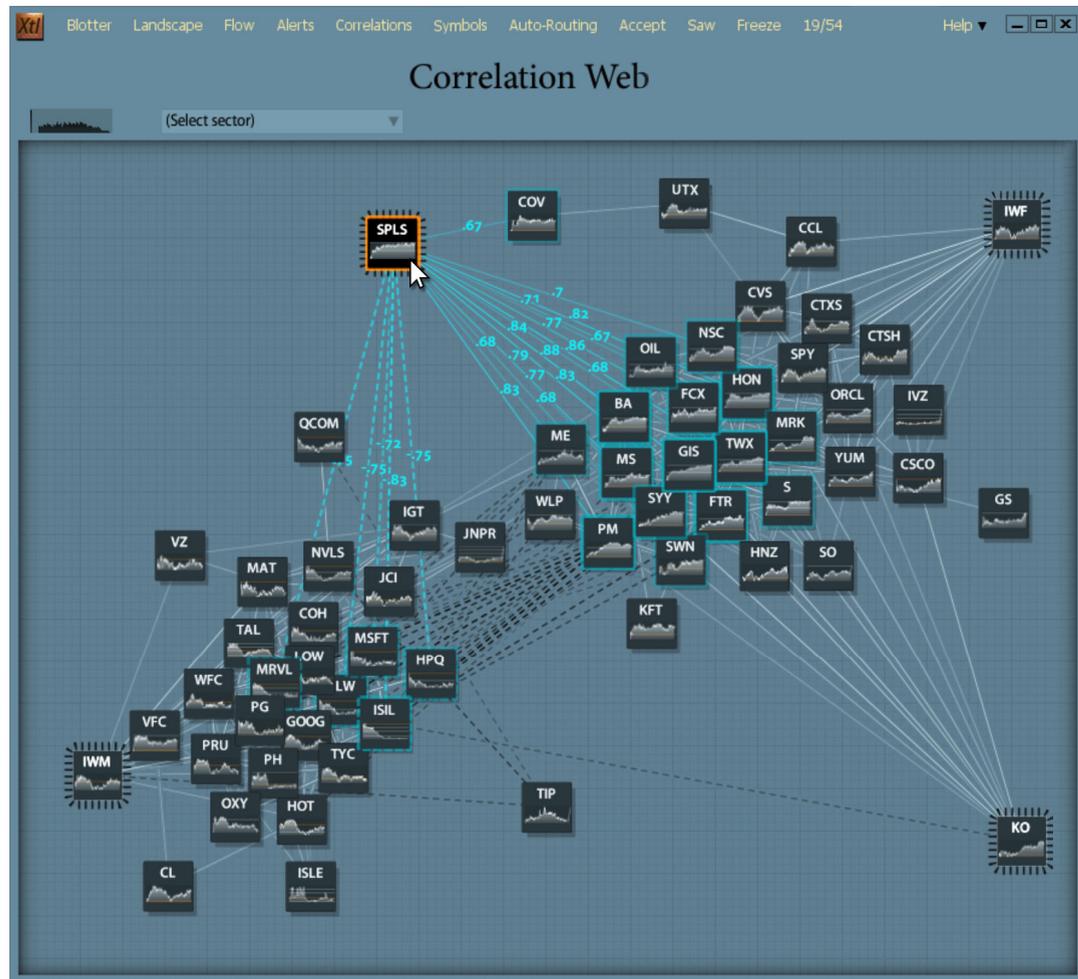
Many graphics present as if they're more important than the viewer's thoughts.



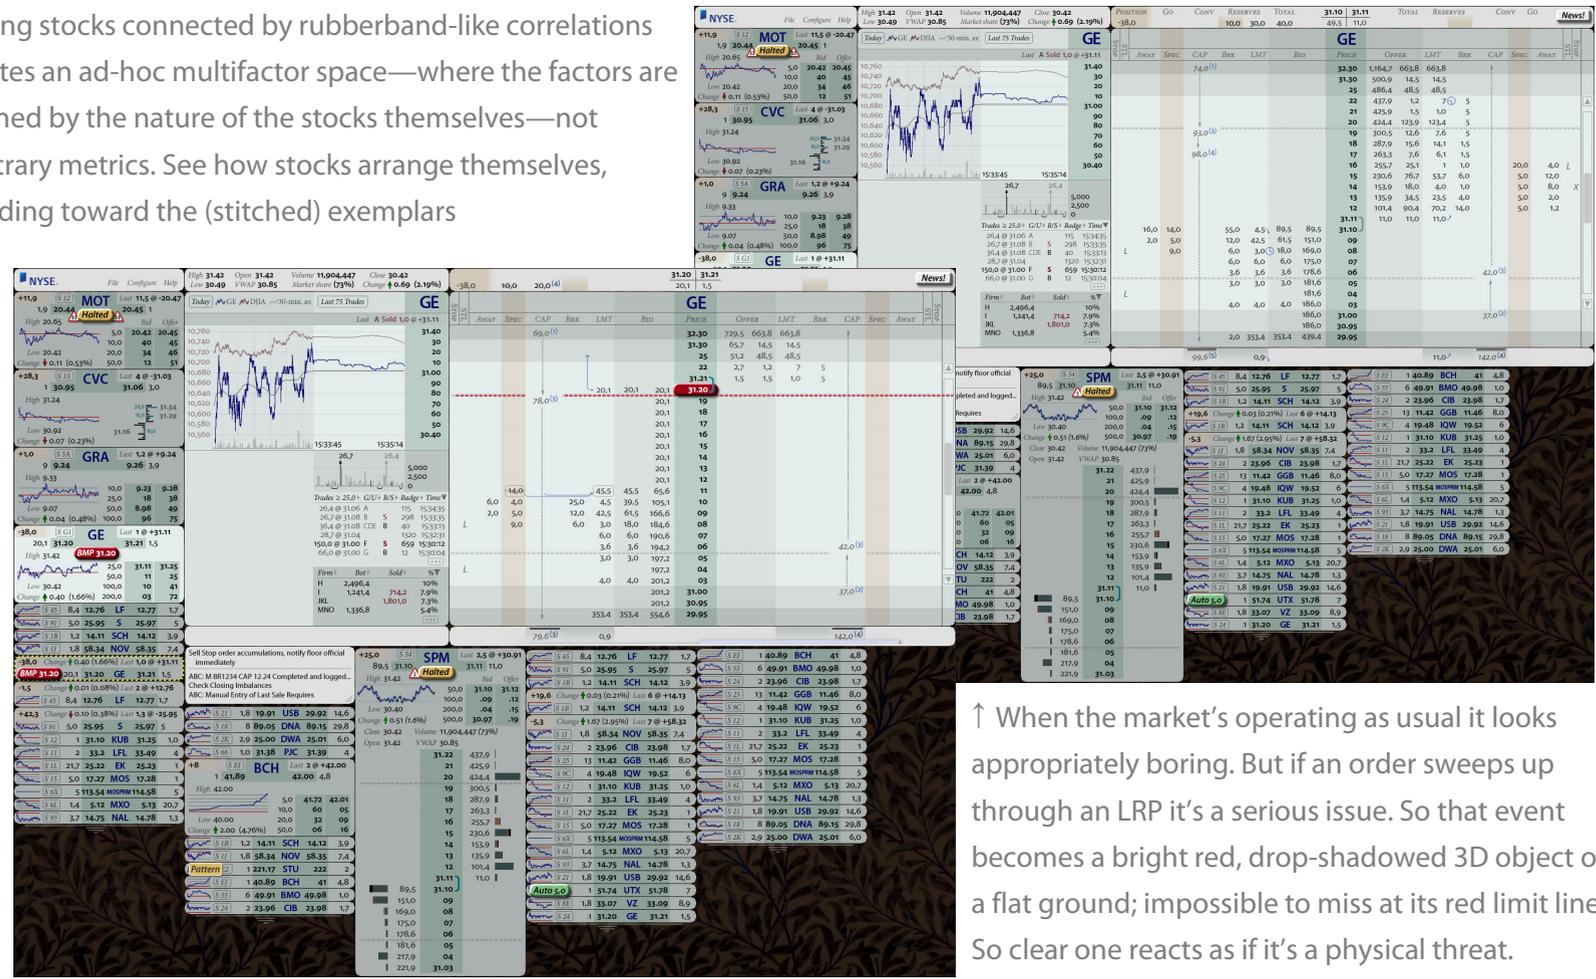
When the game is to grab a website visitor's total attention and keep their eyeballs for as many seconds as possible, it's important to use saturated, contrasty colors and tricky, thought-entraining animations. But real tools—tools glanced at hundreds of times a day—need to be humble visual backgrounds for our thought processes: clear but not demanding, structured to allow easy annotation. True innovation can still be done with basic typography and ancient colors. ("Rubification" of a manuscript was done with inks made of ground rubies.) Left: Once More Around the Sun, a year-at-a-glance calendar; every hour is visible to be marked on the full 2' paper version. Right: a linear calendar that breaks only at months to allow one to arrange blocks of time throughout a year. (Here, my conferences/talks in 2006; a couple had to go—and a little time was needed for work in between...)

Agency Are you immersed in your task—or coping with some tool?

A good tool acts as an extension of your body, your mind. You can even react before you know an issue exists.



Pulling stocks connected by rubberband-like correlations creates an ad-hoc multifactor space—where the factors are defined by the nature of the stocks themselves—not arbitrary metrics. See how stocks arrange themselves, trending toward the (stitched) exemplars



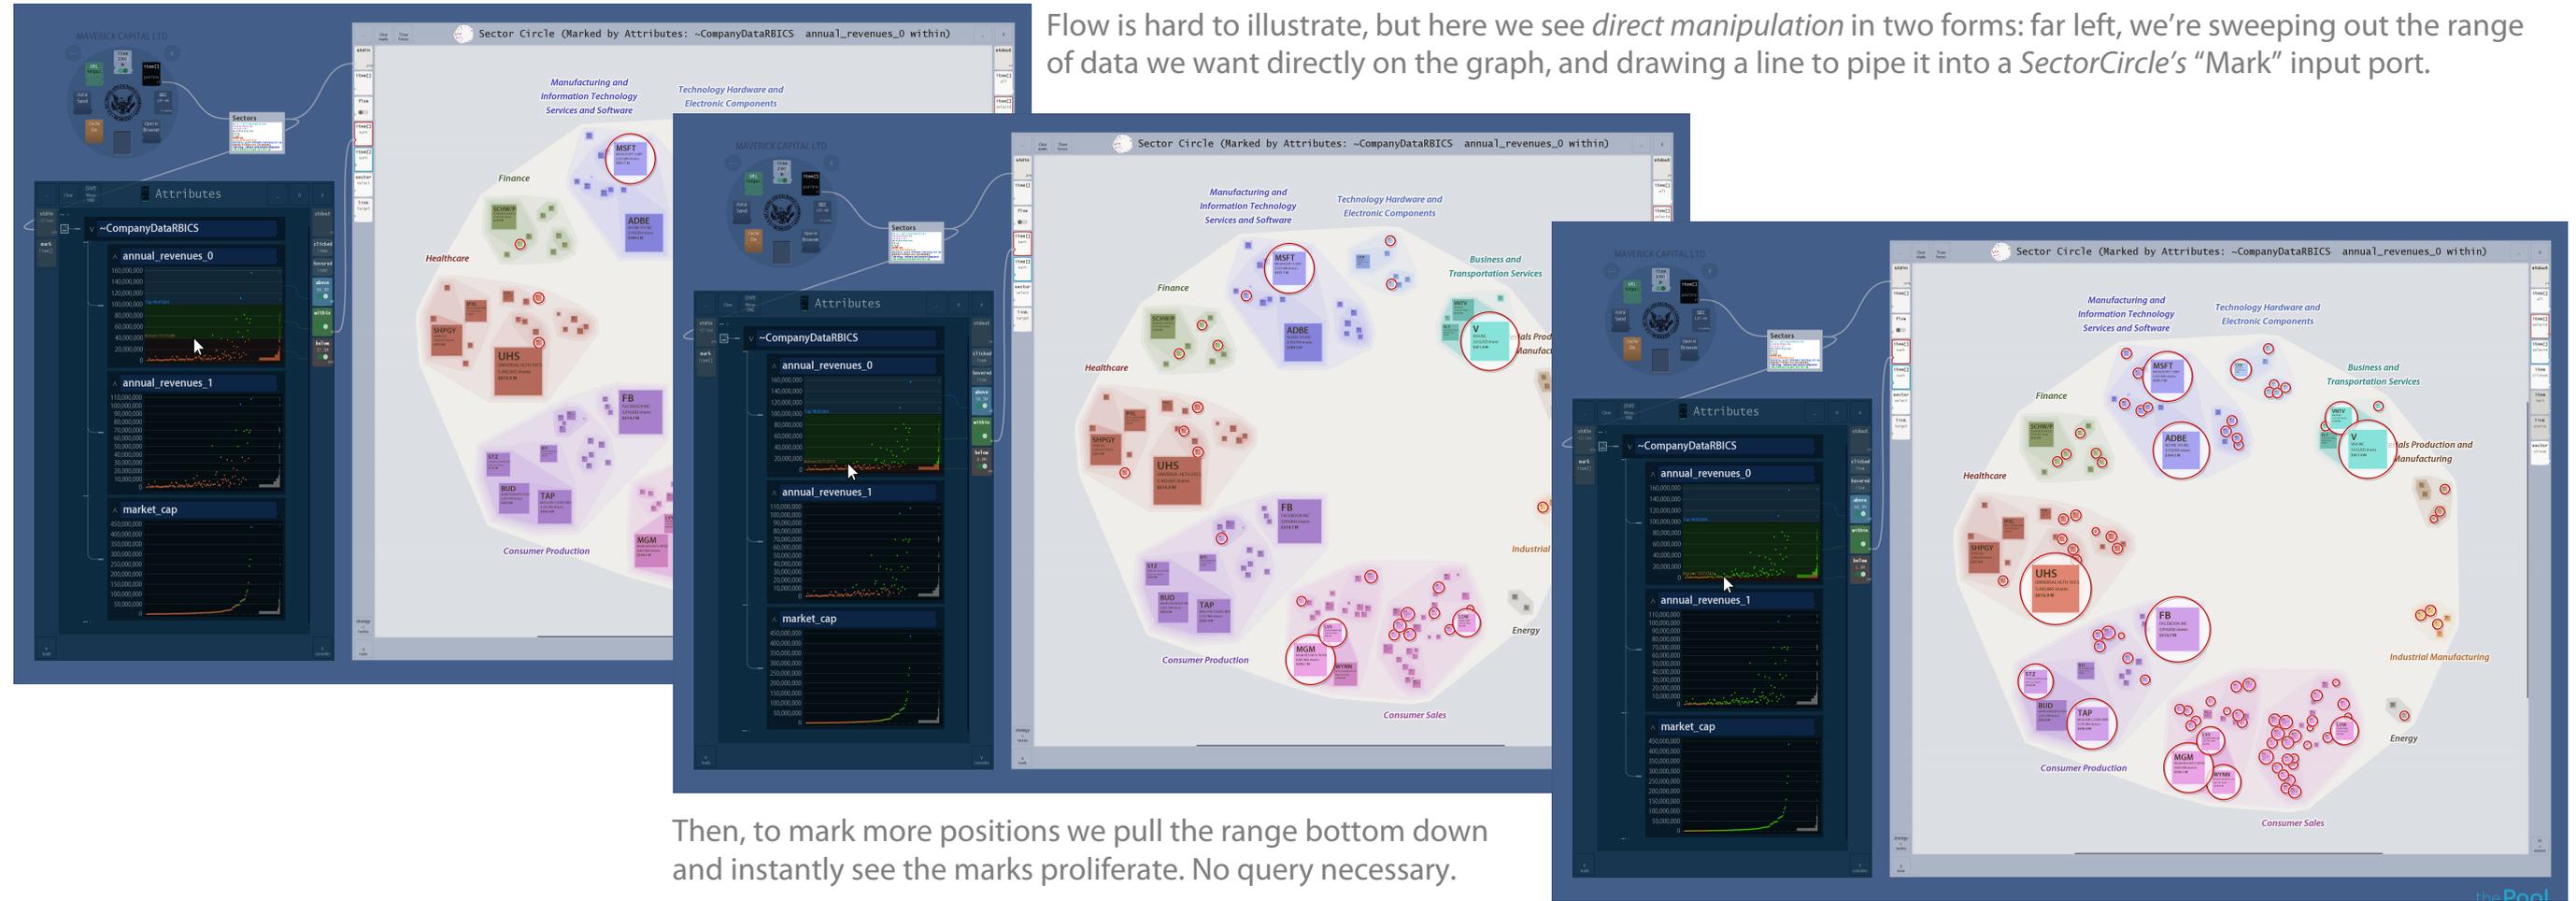
↑ When the market's operating as usual it looks appropriately boring. But if an order sweeps up through an LRP it's a serious issue. So that event becomes a bright red, drop-shadowed 3D object on a flat ground; impossible to miss at its red limit line. So clear one reacts as if it's a physical threat.

Marshal artists have known for centuries that it's possible to react to something before you consciously know it's happening. It was ascribed, e.g., to one's chi being mystically in tune with the external world; Star Wars' the Force. We now know it takes roughly 300ms for an event to enter consciousness but less than 100ms for deeper perceptual mechanisms to recognize certain kinds of things. Related concepts can even be activated in that 200ms—like mental scripts of how those things act and how to respond. Tools that need conscious processing, like words in rectangles, can't tap into this as well as tools that present perceptual features that older mechanisms can process. When we do that there's a sense of agency; it feels like it's us working directly in the world, just doing our task—not like the tool is mediating. This foregrounds tool behavior as essential; people copied our GS NYSE handheld and failed—they copied just the look.

Flow We feel better—and think better—when we’re “in the zone”

Evolution tuned us to see and react in tenths of seconds, so even a half-second refresh hobbles thought.

Flow is hard to illustrate, but here we see *direct manipulation* in two forms: far left, we’re sweeping out the range of data we want directly on the graph, and drawing a line to pipe it into a SectorCircle’s “Mark” input port.



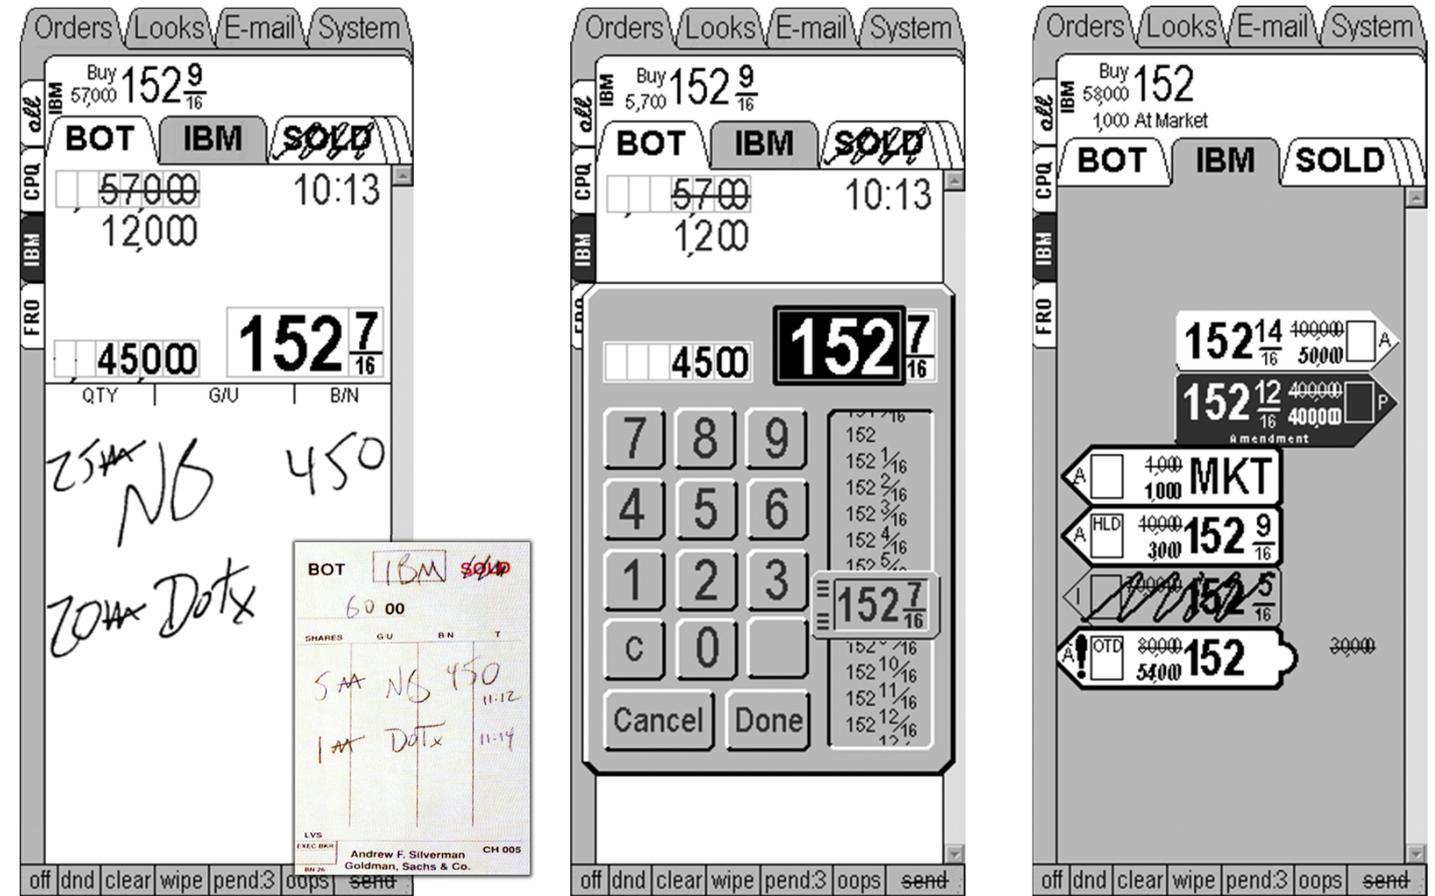
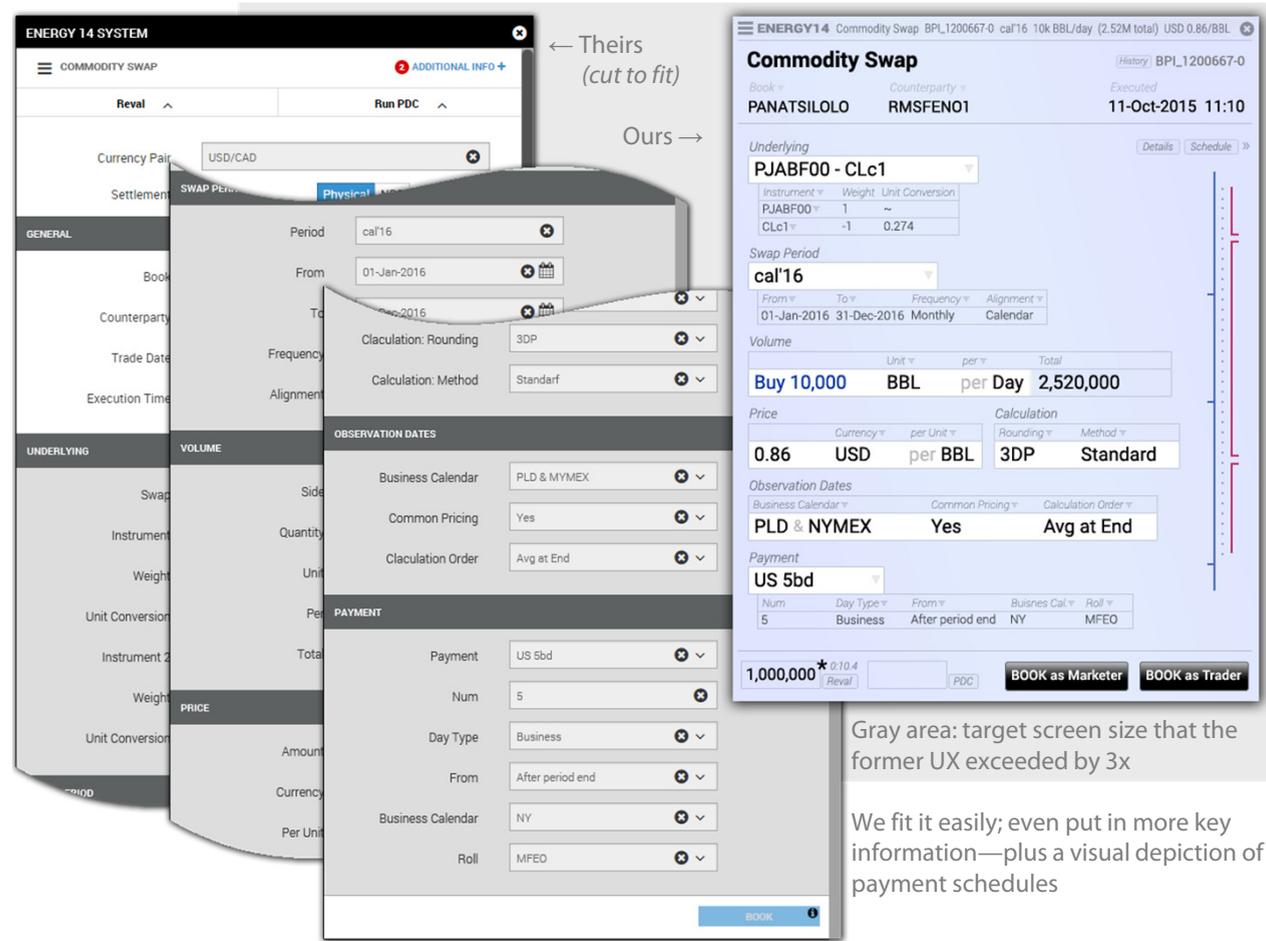
Then, to mark more positions we pull the range bottom down and instantly see the marks proliferate. No query necessary.

Mihaly Csikszentmihalyi studies the psychology of optimal experience. But in the zone your whole-mind behavior goes deeper than the veneer of experience. In fact, flow’s pleasant experience may be a by-product of working effectively in the world: an evolutionary reward making it more likely to be selected into the next generation. Csikszentmihalyi lists conditions that induce flow, and this portfolio shows how we can create those conditions. Notably: a tool that looks like a specific task reinforces a clear goal, and the sense of agency signals the clear, immediate feedback a good tool provides: Ben Shneiderman’s “direct manipulation” with “rapid, incremental, reversable” operations. When a Web tool takes time to respond it’s far worse than a lost ½ second—it interrupts flow, which can take more than 23 minutes to regain. This ruins productivity and alienates the viewer. The Web might someday grow out of its “terrible twos,” interrupting us because it feeds on our attention, or it’s poorly implemented. But a mature set of tools avoids this already, today.

Chunking & Functional Encoding *Manage visual complexity*

Facilitating mental “chunks” lets people juggle more.

Ironically, we often work not to attract attention but *shed* it.

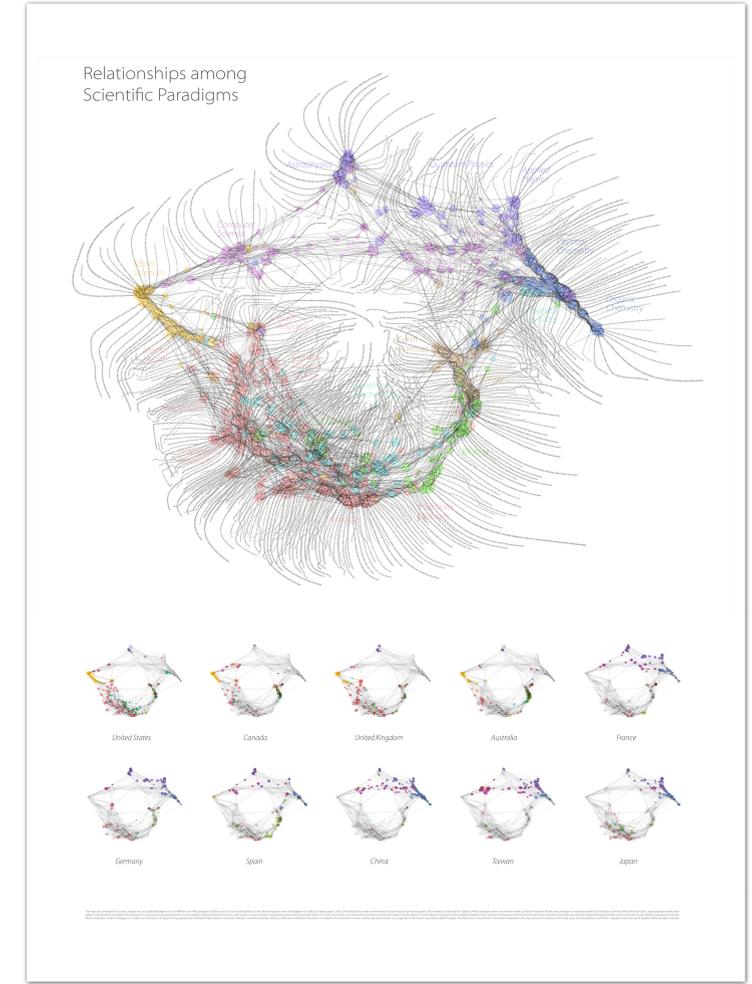
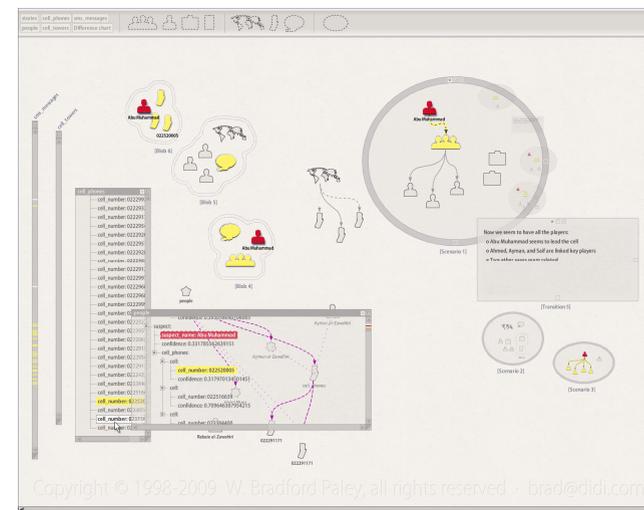
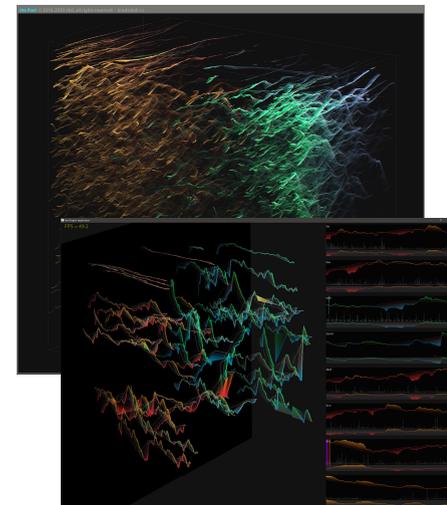
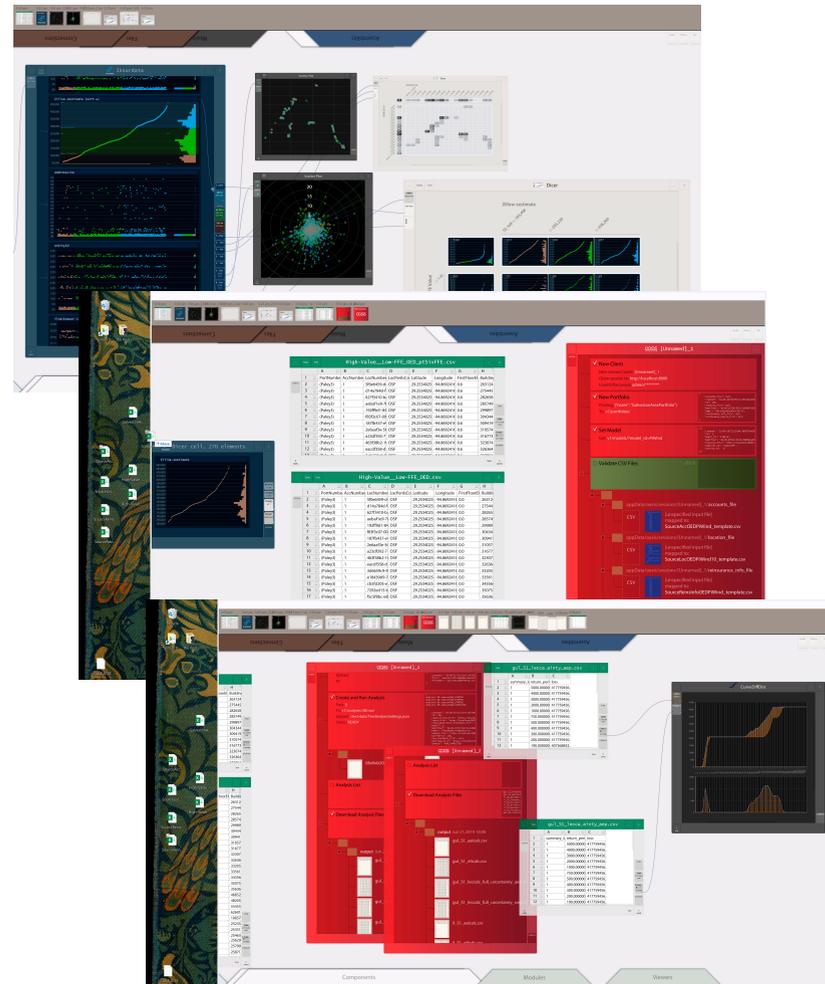


Above: another bit of history, the Goldman broker handheld from 1999, has strongly-differentiated visuals for different functions so your eye snaps to what you need—by letting you ignore the things you don’t, just like in the real world. The fashionista denigration of skeuomorphism misses an essential point: we animals need those rich visual attributes to encode real-world objects by function. Imagine if someone redesigned your home with today’s UX widgets for every function—and there’s a fire. Do you survive? Left: The ink-entry screen copies a real execution slip (inset) because it works—it evolved over decades to an almost ideal size & configuration. The keypad can’t be mistaken for anything but. And when we needed to show several order tickets for brokers to choose among, we transcribed the meaning, not the ticket. Visuals let one instantly see what’s buy or sell, what’s amended/working/done, and importantly: let brokers ignore all of the order details except where sells touch buys at the market. This design kept the Goldman brokers the fastest on the floor, competitive for years after electronic trading started at the NYSE.

We “rescued” the traders at a major international energy-trading bank from a design created by a salient global markets UX design firm. Theirs followed all the rules; our followed the way traders think. Making a separate field for each input certainly created a tidy-looking scrolled strip. But. Broke. The. Thoughts. Of. The. Traders. Into. Separate. Pieces. Worse. Than. This. Does. On the trading desk we just listened. Those pieces were used in sentences—each one a meaningful chunk in their workflow. So we created the Composite Field, and cemented them back together. An unsolicited email from the project director reported: “feedback has been positive, people booking complex products without having to understand the application or been given training.” Customer satisfaction, zero training costs, zero support costs, and no scrolling. All done within a flat design vocabulary.

Co-workspaces & Presentation *Hosting distributed cognition*

Peers think *with us* when visuals act as shared memory, and the ideas are right there for people in other roles.



Edwin Hutchins studied how the five roles on a Navy bridge become a single cognitive process. Their tools (map, compass, log book, even the physical throttle position) bind their thoughts into a single emergent intelligence, smarter than any one person involved. Our tools enable meta-minds to form by optimizing the flow of information between the people involved. Left: the Pool allows global collaborative ad-hoc exploration of any data. Center: Paley is interviewed at the Whitney Museum in front of his CodeProfiles, which reveals coding and optimization processes to a lay audience. Under: the experimental AverageDragons compare hundreds of stocks through time. Next: WorldMerge let the intelligence community synthesize quantitative data with qualitative data. Right: a map of science (featured in Nature and a dozen other publications) presents the whole world as context at the top, so national variations are clearer underneath. Zoom way in to see the West's strength in Medicine, the East's in Math & Physics, and Europe's balance.

Background

W. Bradford Paley *founder, Digital Image Design Incorporated*

Analysis/Transaction Front Ends *(none have more experience innovating here)*

Goldman Sachs NYSE broker handheld; **Merrill Lynch** Equities EMS; **Morgan Stanley** FX EMS for Sales- and Market-facing traders; **NYSE** NYSE's own broker handheld; Specialist station; **REDIPlus** EMS; **BIDS Trading** Dark pool complete front end; **FFI** "TRADEWORKS killer" Fixed Income Broker Station: prototype; **JP Morgan Asset Management** Next-gen OEMS: prototype; **Tora Trading** Complete EMS; **Merrill Lynch** Tora's EMS white-labeled; **Macquarie** 30,000+ (!) Energy Trading screens; **Lehman Brothers** InfoVis for neural networks, other advanced research—in 1989; twice as many non-transaction systems for similar names; **In the works**: a PM/Trader shared workspace, a strat/algo "Clear Box" tester/viewer/controller

Academics and Recognition

Adjunct Associate Professor, Computer Engineering Department, **Columbia University**, (appointed to teach his own work as a graduate seminar); **U.C. Berkeley** Phi Beta Kappa degree in Economics (4.0 in major and minor); **ID Magazine** Design Distinction Award, **Tokyo International Media Festival** Grand Prize [Non-Interactive], **Ars Electronica** production award; **NYSCA** grantee; **NYFA** fellow; **Science Friday** public radio interview

Publication Features, Significant Book Inclusions

The New York Times (full articles on TextArc, CODEDOC), **Nature**, **Geo**, **Slashdot**, **Esoteric.codes**, **Visual Complexity** (Website and book), **Rattle** (featured as visual poetry), **ID Magazine** (feature on GS NYSE handheld), **Atlas of Science**, **Database Aesthetics: Art in the Age of Information Overflow**, **Living With Art**, **Digital Art**, **Wired**

Museums and Galleries

The Museum of Modern Art, **The Whitney Museum of American Art** (permanent collection), **ASU Gallery** (solo show), **The Chelsea Art Museum**, **Streaming Museum**, **Pace Digital Gallery** (the university), **Cooper Hewett**, **Smithsonian Design Museum** (permanent collection, upcoming), People's Republic of China **Central Academy of Fine Arts Museum**, (upcoming)

Keynote/Plenary Addresses

Graph Drawing, Diagrams, 9th International Conference on Information Visualization, **SmartGraphics**, **Astea/FMI** ("Bulgaria's MIT") *2020 Summit*, **Institutional Investor Financial Technology Forum**, **UCLA IPAM Social Networking**, **IEEE InfoVis**, **ALT.CHI**, **The Kitchen** [*& scores of invited talks*]

Contact: brad@didi.co