

Annotated Portfolios

Bill Gaver

Goldsmiths, University of London | w.gaver@gold.ac.uk

John Bowers

Goldsmiths, University of London | john.m.bowers@gmail.com

As soon as a divide is made between theories and what they are theories of, the tip of technoscience is immediately shrouded in fog. Theories, now made abstract and autonomous objects, float like flying saucers above the rest of science, which by contrast becomes ‘experimental’ or ‘empirical.’ —Bruno Latour

A few years ago our team was involved in a project designing for older people living in a care home. Having known the residents for over a year, we appreciated the complex mix of loneliness and sociality, withdrawal and interaction, humor and sadness that characterized their lives. Over that time, we had sketched numerous ideas, produced design workbooks, and even run a daylong session in which they tried a number of lo-fi experience prototypes. Now it was time for us to develop a specific design. What would we make?

You might expect that we turned to theory for direction—after all, our studio routinely produces and engages with theoretical work. And indeed, we did have some conceptual themes we wanted to exercise through the project. Our theoretical convictions, however, turned out to give very little guidance about what, specifically, we should make.

In the end, we developed a trolley called the Photostroller that shows

a continuous slideshow of photographs scraped from the Web, using a controller that enables users to select from a set of predefined categories and tune the degree of “semantic drift” between successive images (Figure 1). We deployed the device in the care home for several months and found that our design supported a range of engagements in the home. Sometimes it fueled conversations in the lounge; other times it was a source of contemplation for a person alone in his room. Overall, the design seemed successful, embodying many of the qualities we wanted.

Still, it is questionable how much anything one might call theory was of direct help in coming up with the design. Instead, we drew from a variety of concrete examples of other work for inspiration. For instance, one of our collaborators, Mark Blythe, had been running Web-browsing sessions with the care-home residents, and they had shown interest in the kinds of photographs found on Flickr. We had also witnessed their enthusiastic reception of photographs we collected to show them in a lo-fi prototyping session we ran. Influence also came from specific designs we knew—such as our own Local Barometer, or the Listening Post by Ben Ruben and Mark

Hanson—that focused on reframing content taken from the Web. And as we finalized the design, we were deeply influenced by Dieter Rams’s design work for Braun.

Theoretical reflection might have been helpful in emphasizing the sorts of qualities we wanted for our design. But our design choices were *underdetermined* by theory—we could have developed any number of things that would have been in accord with our ideas. Instead, it was by looking at specific *examples* of practice that we found guidance for our work and, in discussing exactly how those examples were relevant to us, began to develop our design thinking.

Later, when the field trial was over, we wrote a report of the project. In addition to describing our process, the design, and the results of our field trial, we discussed the project in terms of conceptual themes we thought might be generalizable to other designs. By abstracting the actual features of the device, using terms such as *constraint*, *control*, *drift*, and *openness*, we sought to suggest dimensions of design that might be applied broadly.

However, we could have theorized the Photostroller in any

► Figure 1. The Photostroller with annotations.

influencing
autonomous drift

reframing online
content

constraint and
openness

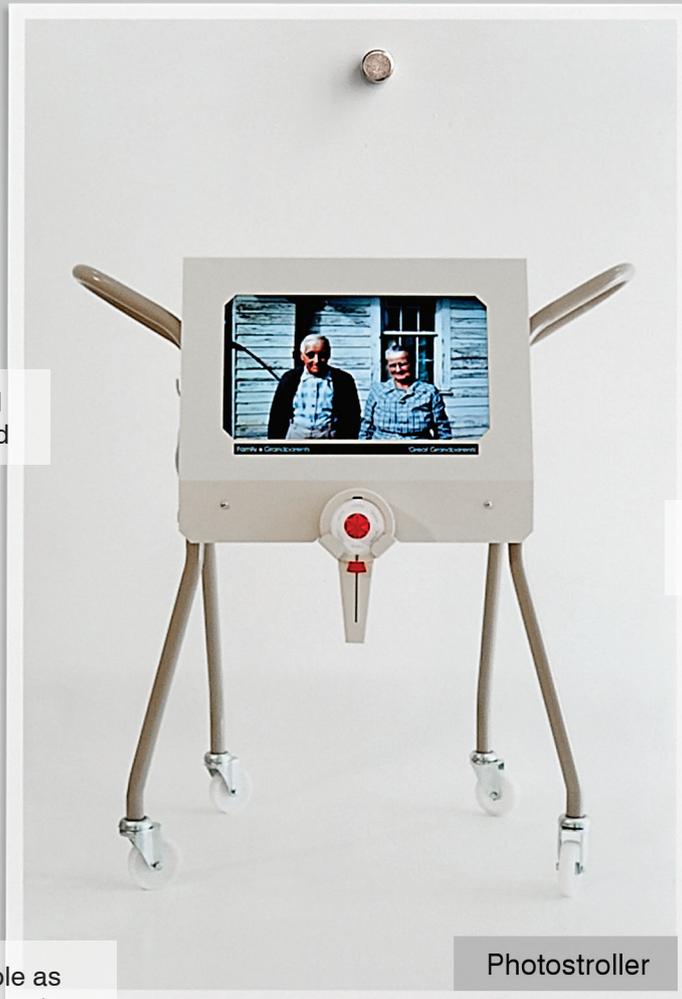
device as threshold
to surrounding world

form design for
everyday settings

framing older people as
curious and engaged

Photostroller

interaction techniques
for variably abled



number of different ways. For instance, we could have focused on general statements about reframing Internet data. We might have centered our discussion on the nature of conversational resources. Or on devices that allow smooth transitions between focused and peripheral attention. Or the design of navigational cues for such an interface, or the way we drew on Dieter Rams's designs in achieving the aesthetic style of the piece, or the health and safety requirements of designing for a care home. Any of these themes could have been abstracted, conceptualized, theorized. Instead, what we highlighted reflected our own views, not only about what was salient for the Photostroller, but also about what is topical and important for research in interaction design.

So, theory might have been helpful in articulating the qualities of our design that we thought important for its success, and particularly those we would like to see extended in the future. But our theorizing was *underdetermined* by design—we could have developed any number of theories that would have fit our design. Moreover, our theorizing was, in a sense, *political*. We did not discuss all aspects of our design and the ways people use it. Instead we focused on those aspects we want to promote in future designs. Our theoretical work didn't just concern what is, but what *ought to be*.

How Can Design Be Research?

In our experience, theory plays an unexpected, indirect, and sometimes rather humble role in design practice. This loose relationship with theory—and often a tighter relationship with other instances of design practice—is not unusual for our work. The question is: Where

does this leave design practice as a form of research? According to some common perceptions of what research is all about, the development of consolidated theoretical knowledge is crucial. So, how could design count as research?

Over the past several years, design-led work has increasingly been accepted as making first-class contributions to HCI research and practice. Coinciding with a general trend for computing to be applied in all aspects of everyday life, design seems to offer the ability to reflect emotional, aesthetic, cultural, and critical concerns alongside those of functionality and usability. The outcomes offered by design often take the form of prototype products and systems, sometimes developed to a high degree of finish technically, physically, and aesthetically, and sometimes accompanied by accounts of field trials of these products in use or in exhibitions. In addition, “manifestos” occasionally appear, arguing for the value of “supple” or “ludic” or “reflective” design as a direction for future work.

It seems clear that the products and concepts of design can reflect insights and innovations that are topical and valuable to the HCI research community, and that moreover have a character that is often distinct from the systems developed by other disciplines. The methods and methodologies that designers use to come up with these products seem, by implication, to have similar value. But is that enough to make design a form of research, or is it merely fodder to be turned into research by others? Do we need to add some special ingredient—predefined research questions, perhaps, or methodological rigor, or the production of theory—to design practice in order

for it to count as research? Or do we have to go further, and change design practices to make our contributions to HCI look a bit more like commonly understood versions of research? In that case, is the result still design, or have we lost something in the process?

These sorts of questions have been vexing the HCI design community—and us—for some time. The problem is that novel products alone do not seem sufficient to count as research, and even a stream of locally innovative designs may not seem to add up to much. But strategies for generalizing beyond practice to produce a body of work recognizable as research seem equally fraught with difficulty. Methodological frameworks promise rigor but jeopardize the possibility for designers to invent ad hoc approaches, or draw inspiration from unorthodox sources, or take inexplicable imaginative leaps—all forms of a productive indiscipline that we see as integral to design practice. Theory promises generality and guidance but seems inadequate to capture the situated, multidimensional, and configurational nature of design, and moreover threatens to occlude the potency of unique, embodied artifacts in a cloud of words and diagrams. Methodologies and theories may well produce respectable research, but the danger is that this will come at the expense of design.

We believe a closer look at endemic design practices might provide a clue about how best to manage the tension between theory and practice in design research. Rather than importing tactics from other disciplines to fashion new normative guidelines, we suggest the answer may be to recognize and build upon practices already in use within the

field. This would allow designers to legitimate their activities as research, not by changing their practices to resemble those of other fields, but by more careful articulation of their existing approaches in a way that is legible to those from other disciplines.

In particular, we suggest that *annotated portfolios* provide a way to present the fruits of design that simultaneously respect the particularity and multidimensionality of design work while meeting many of the demands of generalizable theory. In the rest of this article, we describe the logic behind this notion, illustrate it with an example of an annotated portfolio of our studio's own work, and describe some of the benefits we see for this strategy of simultaneously illuminating the material and conceptual dimensions of design research.

The Logic of Annotated Portfolios

A myriad of choices go into the design of any given artifact. These choices are varied, multifaceted, and heterogeneous. They reflect a very wide range of concerns that may include:

- the *functionality* of the design (what should it do?) and by implication the value of certain activities (is this worth doing?)
- its *aesthetics* (what form and appearance should the artifact take?)
- the *practicalities* of its production (what materials, skills, and tools are needed to make it?)
- the *motivation* for making (why are we doing this? what are we trying to show?)
- the identities and capabilities of the *people* for whom the artifact is intended (what will our users make of this? how can we best design for them?)
- *sociopolitical concerns* (what sort

of culture will this encourage or resist?).

From this point of view, a designed artifact can be seen as a kind of position statement from its designers, not only about what is important to consider in a given design situation, but also about how to best respond to those considerations.

The trouble with this perspective on artifacts, however, is that neither dimensions of concern nor designers' orientations to them can be read directly from the artifacts themselves. Any given design can be viewed from an indefinitely large number of perspectives, only a few of which may be salient to the designers themselves. This doesn't mean that other views are inaccurate or irrelevant, of course—for instance, designers may well be unconscious of a host of cultural, economic, and political assumptions embodied in their designs—but it does render problematic the notion that artifacts embed designers' research in a way that can be appreciated directly.

Instead, as designers we typically comment on the artifacts we build to highlight the concerns and choices made in their development. This happens routinely both in commercial and research settings. We advertise the new and improved features of a product, for example, or describe how a prototype has relevance for a hot research topic. In both cases, we point out what makes the design new and valuable, rather than leaving the artifact to speak for itself—as if it could.

Given the multifaceted considerations that go into any design, however, it is impossible to describe a given artifact in all its detail, from the political philosophy that informs it to the tactile quality of a given toggle switch. Instead, our

written accounts can at best be partial views onto the design as a whole, focusing on some concerns at the expense of others. They are also likely to be incomplete in another respect: Much of our knowledge of making is tacit. As such, it need not appear in written texts because we assume it is shared by members of the community or we will be contacted directly by those seeking clarification.

This means that textual accounts (e.g., published papers, catalog entries, online descriptions) in design research have an *indexical character*. That is, they point to features of our designs and connect them to matters of further concern, in the case of research, making them topical for discussion within a given community.

From this point of view, artifacts do not merely demonstrate assertions made more precisely in a theoretical account. On the contrary, we see textual accounts of artifacts, including any theoretical pronouncements about them, as *annotations*. The textual account achieves its sense and relevance by virtue of its indexical connection with an artifact. At the same time, aspects of that artifact are highlighted and linked to the concerns of a community by the text. Just as a pointing gesture depends on both the finger and the target, and as a caption makes sense of, and exists because of, its associated figure, so artifacts and their descriptions are mutually reliant on their relationship to produce meaning.

This line of reasoning implies that designs need to be annotated if they are to make clear and accountable contributions to research. Such annotations not only point to salient features of an artifact but also bridge between the artifact and issues of concern to the research

reframing online
content

influencing
autonomous drift

device as threshold
to surrounding world



framing older people as
curious and engaged



► Figure 2. The Photostroller juxtaposed with the Prayer Companion brings breadth and depth to shared annotations.

community. Insofar as such issues include conceptual ones, annotated artifacts can be seen to serve some of the same roles as theoretical statements, identifying activities and qualities deemed important for research and suggesting ways these might be addressed.

Beyond single artifacts, however, annotated portfolios may serve a much more valuable role as an alternative to more formalized theory in conceptual development and practical guidance for design. If a single design occupies a point in design space, a collection of designs by the same or related designers establishes an area in that space. A single artifact embodies propositions about a specific configuration of properties. A comparison of multiple items in a portfolio, on the other hand, can make clear a domain of design, its relevant

dimensions, and the designer's opinion about the fruitful locations and configurations to develop on those dimensions.

Building a Sample Portfolio

An annotated portfolio, then, is a means for explicating design thinking that retains an intimate indexical connection with artifacts themselves while addressing broader concerns in the research community. To illustrate this notion in more depth, we consider examples from the body of work we have been developing at the Interaction Research Studio over the past 10 years or so.

Let us start with a single artifact—the Photostroller we described at the beginning of this article. In Figure 1 we label an image of the Photostroller with a number of very brief annotations, annota-

tions that gain their sense from the Photostroller itself and reach out to broader matters of interest in our research communities.

For instance, consider the annotation *interaction techniques for the variously abled*. Taken by itself, this could refer to any of a huge number of techniques and underlying approaches, from the employment of extra-large controls or augmentation of familiar objects to the provision of specialized feature sets or social networking support. Juxtaposing the annotation with the Photostroller, however, makes clear that the annotation refers to the specific technique being proposed by this research—the particular combination of category selection and specification of “semantic drift” offered by the device's controller. At the same time, the annotation links this control scheme explicitly to

constraint and openness

form design for everyday settings

interaction techniques for variably abled

the class of techniques for improving interaction for the variously abled, allowing its contribution to be compared with others. In this way, the annotations build bridges between larger research issues and the approaches to these issues embodied by the PhotoStroller itself. This linkage does not, however, come through abstracting away from the artifact. Instead, our approach to larger research issues is concretely exemplified in the PhotoStroller—and with the precision that building a highly finished prototype requires, and thus offers when it is complete.

Now let us add another artifact to this picture, the Prayer Companion (Figure 2). This was developed as a resource for the spiritual activity of a group of cloistered nuns. It displays a stream of information sourced from news feeds and social

networking sites to suggest possible topics for prayers. It was developed in the same overall project as the PhotoStroller, one devoted to exploring technologies for older people. Not surprisingly, the two artifacts share a number of features. The juxtaposition of the two artifacts, however, gives breadth as well as depth to some of those annotations and motivates comparative discussion over others.

For example, both the PhotoStroller and the Prayer Companion construe their senior users not as individuals requiring medical care or assistance with living, but as people who are actively curious and engaged with the wider world. Both devices manifest an interesting interplay between constraint and openness in design. On the one hand, they are quite narrow in their interactivity and functionality. The PhotoStroller “just” shows images and can be influenced with only two controls (category selection and drift control). The Prayer Companion “just” shows scrolling topical text in a manner not controllable by the end users at all. However, this simplicity of functionality leaves a great deal open-ended. Exactly what people will make of the images or text items is an open affair. Just how the devices will embed within everyday contexts is left for users to determine. Though we anticipated the devices resourcing conversations and other forms of social exchange (including prayer) important to our users, exactly what content these interactions take is not mandated by us.

The PhotoStroller and the Prayer Companion were both designed to have autonomous “drifting” behaviors. They will continue to stream their content without the intervention of a user. This content will also periodically change as

new images or textual items are retrieved from the Internet and added into the pool. In the case of the PhotoStroller, this process can be influenced through the operation of the wireless controller. While the Prayer Companion varies the content of its display and other features, such as scrolling speed and degree of repetition, to pique visual and intellectual interest, it does this autonomously, without user control. By comparing devices with respect to the annotation, we can begin both to capture the “family resemblances” that exist in the emerging portfolio, and to see how the designs embody a range of different but equally valid (in our studio’s view) responses to the topics indicated by the annotations.

Let us now add three further artifacts (Figure 3). The Drift Table displays slowly moving aerial photography controlled by the distribution of weight on its surface. The Local Barometer displays online text and images related to the home’s locality, depending on the local wind conditions to give an impression of the sociocultural surroundings. The Plane Tracker tracks aircraft passing overhead and imagines their flights onscreen to resource an understanding of the home’s global links. In different ways, all of these involve the reframing of online content and serve as a kind of threshold between the locality of their use and the wider world. In each case also, the form of the device has been carefully crafted to be mindful of several concerns: the everyday settings in which it is likely to be used, the affordances of the materials and technologies used in construction, culturally significant aesthetic traditions that are drawn upon, and so forth.

It is important to recognize that annotations and the designs they

annotate are mutually informing. Annotations have an indexical relationship to the artifacts they are relevant to. Annotations depend on traceable connections to design for their significance, just as designs are illuminated through annotation. For example, the annotation *device as threshold to the surrounding world* only really gains its sense if one looks at the artifacts in the portfolio it is true of—nearly all of them in one way or another. However, the Drift Table is true to this in a different way from the Prayer Companion, and one needs to appreciate the particularity of those designs to understand the specific sense of the annotation that is relevant.

Useful annotations, then, tend to have a *general application* in the portfolio yet a *specific sense* for individual artifacts within it. Moreover, the sense of an annotation is enriched by reference to example artifacts: It is not just broadened, but given shape and dimensionality by the examples. Individual designs illuminate ways to address the issues raised by annotations and, when accumulated, suggest patterns of similarities and differences in those strategies. They configure a “design space,” a zone of potentially fertile possibilities.

Let us further elaborate the design space in which these artifacts exist. Figure 4 focuses on users’ experiences with the designs: how users have appropriated the devices in actual use, the senses they have found in them, and the everyday activities in which they are embedded. A key observation we have made a number of times is that our designs *resource social interaction*. They serve as conversation pieces or curiosities, to be sure, but, more than that, the content they deliver and how they do it typi-

cally has a balance of *openness* and *constraint* that facilitates interaction between people over matters of interest and concern to them. The Photostroller, for example, enabled the older people for whom it was designed to reminisce, through the images it showed, about their locality and how it had changed in their lifetimes. It also created the opportunity for discussions in which people compared their likings and dislikings of what they were seeing, and, broadly speaking, their appreciation for photography itself.

We have observed our devices being used in a variety of ways, many of which combine *playfulness* and *provisional* tasks. For example, we designed the Drift Table to support a kind of curiosity-driven play, rather than particular functions better served by more goal-directed tools and interaction techniques. However, our design does not prohibit people from setting provisional tasks for themselves, in which a more concrete aim is locally formulated, for example (to cite an example we observed), using the Drift Table to travel to Bournemouth to look from above at the homes of the rich. Accordingly, we annotate Figure 4 to highlight this mixture of play and situated task-orientation that we commonly observe in how our devices are encountered in practice.

Features of Annotated Portfolios

Clearly, the annotations of our portfolio shown in the figures here are extremely succinct, even when considered along with our elaborations in the body of the article. For some purposes, such brevity might be appropriate, while for others (including presentation to a typical academic community), we might well expand each annotation significantly. Nonetheless, even this brief illustration enables us to draw

constraint and
openness

device as threshold
to surrounding world



out some characteristic features of annotated portfolios:

- Annotations make a collection of designed artifacts into a portfolio. They bring together individual artifacts as a systematic body of work.
- Annotations capture family resemblances between designs in a mesh of similarities and differences.
- Typically a portfolio can be annotated in several different ways, reflecting different purposes and interests and with different audiences in mind.
- Annotations and the designs they annotate are mutually informing. Artifacts are illuminated by

reframing online content



Photostroller



Prayer Companion



Plane Tracker



Local Barometer



Drift Table

influencing autonomous drift

form design for everyday settings

annotations. Annotations are illustrated by artifacts.

We believe that assembling a body of work into an annotated portfolio can have a number of benefits. First, it can give an accessible account of the design style of a particular project, studio, or institution. We feel the annotations we have given don't merely structure the presentation of the devices but go some way to uncovering their underlying aesthetic and socio-political values. As such, annotated portfolios enable us to compare the work of different projects, studios, or institutions. For example, the portfolio of another group might

reveal a tendency to illuminate socio-political issues by eliciting ambivalent reactions to seductively dystopian scenarios. This might be contrasted with our work, which, we believe, avoids both dystopian and celebratory reactions to prevailing trends in favor of providing alternatives grounded in people's specific experiences and identities.

Annotating a portfolio of work can also help us bring out features that may not have been covered in the individual papers that publish accounts, typically, of single artifacts. This can facilitate assessing the value of the scientific and aesthetic contribution of a body

of work. For example, the style of design we have analyzed above has a clear identity and one different from orientations to HCI such as software agents, direct manipulation, ambient computing, emotional computing, and other tendencies. To add to our list of characteristic features of annotated portfolios:

- Annotations communicate the nature of the portfolio and enable its comparison with others.
- Annotations can shape how artifacts are appreciated and understood, and what scientific and aesthetic value they might have, as well as suggest future research and design possibilities.

► Figure 3. Adding three additional artifacts highlights family resemblances, while retaining the particularities of the designs.

constraint and openness

resourcing social interaction



passive and active engagement

play and provisional tasks

► Figure 4. Annotations can address different perspectives on portfolios. Here we focus on users' experiences with the designs.

Finally, in annotating our studio's work here, we have used a mixture of text and images appropriate for a magazine article. Elsewhere we have presented the notion, and our work, in academic papers. A set of annotated photographs edited into a monograph with an introductory essay might situate the portfolio more appropriately for some design audiences. However, it should not be thought that written texts with illustrations are the only way to present an annotated portfolio. Much depends on the audience that one is trying to reach. Many exhibitions—including those drawing together work that might otherwise be seen as unrelated—can be seen as a kind of annotated portfolio. Annotated portfolios might take the form of videos, or a stage show, or a collection of postcards. Thus, any material form and format can be

considered for an annotated portfolio, including an illustrated monograph, a scientific paper, a curated exhibition, and so forth.

Design and Theory

We propose the notion of annotated portfolios as a way to communicate design research. In part, we do this to provide an alternative to accounts that suggest for design to become productive as research, it should engage in some sort of theory formation. While what exactly is meant by *theory* is not always clear, writers usually have in mind some conceptual machinery that can explain and predict. Experiments test theories. While we do not dispute that *sometimes* a designed artifact and the experience of its use can have an experimental status, we do not believe this is typically the case in the kind of design

research we do in HCI. Accordingly, it is moot whether we want to regard annotated portfolios as a contribution to design theory in HCI, if by *theory* is meant something with explanatory and predictive power. Rather than explain existing devices (and this already seems like an odd formulation), we seek to describe them in a way that has an indexical, mutually informing relationship to the devices themselves. Rather than predict the future, we seek to inspire novel work and offer a mapping of the dimensions of emerging design spaces in which it might be situated.

Any particular set of annotations is perspectival, allowing other annotations to be made. Annotations allow family resemblances to be reasoned about, rather than theoretical deductions to be made. Annotations help

assembling resources for engagement and interpretation



Drift Table

us understand what has made a body of work characterful. This may help us understand its successes and failings and inspire future work, but the logic seems to us rather different from that governing theory construction and hypothesis testing, at least as those processes are typically described by writers who call for more rigor in design research or for theoretical or methodological integration with more traditional approaches.

Nonetheless, annotated portfolios may well do some of the work that “theory” is often invoked to do, including addressing the criteria of rigor and relevance. One does not necessarily need the concepts of an overarching theory to be able to describe the world with clarity and rigor, nor is a (explanatory, predictive) theory necessarily required to make one’s work

relevant to the concerns of users or researchers or anyone else. Similarly, it should not be thought that being loyal to the particularities of artifacts necessarily means we have to make enemies of generality. Annotated portfolios are, perhaps, a way of modestly and speculatively reaching out beyond the particular without losing grounding—and doing this with all the rigor and relevance needed to inform the invention and detailed development of new designs.

Part of our enthusiasm for proposing annotated portfolios as an approach to communicating design research lies in the fact that they are not a big step from what designers do anyway. It is common practice for designers (and artists, for that matter) to gather their work into a portfolio and highlight important features that run through it. The presentation of a portfolio will often point to family relationships: “In this piece, I sought to develop the idea from that piece but in a different way.” Or, “These five pieces are all explorations of similar concerns.” And so forth. We feel reasoning about portfolios is a practice that is indigenous to design and, accordingly, many designers in HCI will feel more comfortable working up annotated portfolios than having to integrate their work with theoretical constructs that may not have had a clear role in motivating what they do.

Equally, we are happy to be offering an approach that is open to interpretation and appropriation, in keeping with our own design work. Annotated portfolios do not propose a format of presentation or a set of concerns to be addressed. They do not mandate a particular graphical style, or prescribe a series of categories to be employed,

or advocate an elaborate ontology of entities and relationships. The notion does demand that particular designs be linked to others by way of issues of relevance to the research community, but what those issues are or how they should be expressed is not our concern. In some sense, what we are offering here is a methodology for communicating design research, but not a restricted toolkit of methods.

We believe design research has a characterful contribution to make to HCI that may be watered down if excessively “scientific” criteria for the validity of research are laid at its door. To be sure, ongoing work is needed to maintain and improve the quality of research undertaken through design. Equally, we need to simultaneously bridge to other disciplines while drawing them in to appreciate how design works. Rather than doing this by mutating design to become more like “real” research, however, we believe it is better to grow design’s identity as research from its existing practices and reasoning. Our hope is that annotated portfolios will be one step along that path.



ABOUT THE AUTHORS

Bill Gaver pursues practice-based research on computational products for everyday life. With his team at the Interaction Research Studio, he has made a number of conceptual and methodological contributions through the development and deployment of highly finished prototypes. His current research focuses on batch producing devices that address our relationship with the environment.



John Bowers is concerned with the design and experience of innovative technologies, particularly in public spaces and domestic settings. He pursues a hybrid practice combining software development, empirical social scientific studies, and reflection on design methods. He also plays improvised electro-acoustic music and creates intermedia performances and installations.