

Ergonomics for Interaction Designers

Understanding and applying physical fit in user interface research & design

Rob Tannen, IDSA, Bresslergroup

Driving Factors

Ergonomics is necessary for 3-dimensional, tangible product design where issues of physical fit and comfort are critical. But for interaction designers in the 2-dimensional world of the display screen, understanding physical ergonomics has largely been irrelevant. For example in most cases, interfaces are designed for existing, defined hardware that are out of the control of the interaction designer.

But the continuing convergence of digital interfaces with physical products is putting interaction designers in a position where knowledge of anthropometrics, kinesthetics, and other non-cognitive human capabilities is valuable for creating effective design solutions.

There are several trends contributing to this, including:

1. The rapid proliferation of touch screen and other gestural interfaces which combine "direct" physical control with digital interface design. If you want to design for a finger, you have to know how a finger works.
2. The growth of ubiquitous computing leading to an increased range of scale and form factor in devices that contain interfaces, from traditional computers and laptops, to kiosks, tablets, phones, interactive video walls, electronic ink and consumer appliances (to name a few). As a result, people are interacting with interfaces in range of positions and contexts that go beyond simply standing or sitting in front of a screen. So beyond fingertips, knowing how people can reasonably use their bodies to hold, view, reach and interact is valuable.
3. Computing power and bandwidth across such devices now supports more complex, involved tasks such as data entry, long duration reading and gaming, all of which can lead to risks for repetitive motion injuries, or at least discomfort. Having a knowledge of the types of interactions that can cause such injuries, and how to design around them, is essential.
4. An ever increasingly diverse range of end-users are gaining access to interactive devices, across age, and physical characteristics. For example, the One Laptop Per Child campaign has produced a global, kid-sized laptop. In home health care, a market of predominately elderly users, more devices contain embedded interfaces. And ADA and similar legislation requires that devices are accessible to users with a range of disabilities. In other words, you need to know your user, for it is not you - a given in interface design, a necessity in ergonomic design.
5. Last, but not least - interest. Several of the factors described above are driving many interaction designers to explore and study the world of physical product design. For example, the IIT Institute of Design recently hosted a "thinkering" workshop specifically to provide *"an opportunity for interaction designers to get their hands dirty with electronics, soldering, and wiring, and learn how to interface hardware artifacts with virtual interactions."* Just as it is important to understand the electro-mechanics of hardware, it is essential to understand the relevant mechanical attributes for the users of such hardware.

What all of these trends have in common is a growing need to accommodate human physical characteristics and constraints in the design of digital interfaces. For the most part, this skill set is not part of the experience of interaction designers.

Anthropometrics: The Building Blocks of Ergonomic Design

In any field of design there are those elements that are defined and unchangeable, and those that are malleable. It is the latter in which designers specialize. For example, in interaction design, the fixed elements might include a specified screen resolution, development language and minimum type size. As you might guess, in physical product design, there are many constraints, but human physical characteristics are the most fundamental. Therefore, the most fundamental design question is, how do I design for the range of human physical constraints? For this, we turn to anthropometrics the measure of human body size and proportions.

Suppose we are designing an interactive touch screen kiosk that will be used in an international airport terminal. It is expected that the kiosk users will include travelers from around the world, male and female, from kids through elderly adults. While this may sound like the worst case scenario for physical design (and it is), it's also very typical. In this case we are going to focus initially on *eye height* because we want to set the display so that it can be viewed most easily without looking up or bending down too much.

If we refer to anthropometric data tables, like those found in Stephen Pheasant's *Bodyspace* (2005), we find quite a range in eye height, varying by nationality, age and sex. For example an average, 50th percentile Dutch man has an eye height of 1670mm, while an average, 50th percentile eight year old British girl has an eye height of 1165mm. That's over a 500mm difference, and those aren't even the most disparate populations! So how do we accommodate the diversity of physical characteristics?

Molenbroek and de Bruin¹ discuss the various approaches that one can take to accommodating the range of anthropometric characteristics. The most basic approach, if we can even call it that, is "*Procrustus*", which means that no attempt to accommodate the user has been made, and the user must adapt to the product, however it happened to be designed. Incidentally, this term comes from Greek Mythology, where Procrustes was fitted to a bed by sawing off his head and feet. Only slightly better is the Ego-design approach, where the designer uses his or her own body as a reference. Now every designer does this to some extent for convenience, but it should serve only as a starting reference point.

Design for the mean sounds like a good idea - find the average eye height, and the majority of users will be accommodated. False assumption – in fact a majority of people are excluded by relying on the mean, with only a few falling into the sweet spot in the center.

Designing for one end of the spectrum (small) or the other (tall), can work in some cases. For example, if you design a door to accommodate the tallest users, then by definition, those of shorter stature will fit as well, as clearance is a one-ended variable. But in our case, the appropriate height of a kiosk display is a two-ended issue - there is a hypothetical "too high" as well as a "too low".

Which brings us to some practically workable approaches. *Design for adjustability* means that the product can accommodate a range of users, typically through a mechanical solution. For example, a tilting, height adjustable screen, or multiple interaction stations set at different eye heights. Of course adjustability in the physical world adds cost and complexity, and can lead to unreliable products, so is not always an available solution.

In the end, the most common solution is to *Design for More Types*. In practice this typically means defining a population and then fitting for a reasonable range within that population. Traditionally that range spans from the smallest fifth percentile to the largest 95th percentile. This

¹ <http://www2.io.tudelft.nl/research/ergonomics/AED/publications/enhancing%20anthropometry.pdf>


includes a very broad range of users, but purposely excludes the most extreme 10% of the population (the largest 5% and smallest 5%) - the long tail, where a small number of outlier users can account for a significant design change.

Last, but not least is the ideal - *Design for All*. This means that the product can fit the entire range of an anthropometric characteristic. This is technically possible as humans are not infinitely variable in any dimension.

But having the display at an appropriate height for visibility is just addressing one aspect of interaction - the user also needs to control the interface - in this case via a touch screen.

Designing for Multiple Anthropometric Dimensions

There are several body measurements that could be relevant for reaching a touch screen, but a practical one would be Forward Grip Reach distance - roughly the distance from the shoulder axis to the palm of the hand. With those two metrics in mind - eye height and forward grip reach - you could picture any user as the function of two perpendicular lines - a vertical line, representing the individual's eye height, and a horizontal line representing arm reach. This is illustrated in the accompanying figure for three different representative users - note that the wheelchair user has a sitting eye height compared with the two standing users.



While it might seem relatively straightforward as to how to situate the kiosk - place the screen at a distance and height that accommodates the greatest range of users - the story gets more complicated. The factor that adds complexity is the lack of correlation among anthropometric measurements within people.

In interface design, one is typically working within the constraints of a display. For example, a common resolution for web browsers is 1024 pixels x 768 pixels. Some older displays might be set at 800x600. So while the specific vertical and horizontal dimensions change, the relationship between height and width, or aspect ratio, remains constant at approximately 1.3 in both cases.

So if you're taking a design originally intended for 1024x768 and then need to scale it down to 800x600, it will need to be reduced proportionally.

Ergonomic design would be much easier if people had consistent "aspect ratios", but our body measurements are not predictably proportional or strongly correlated. Meaning that all of the tallest people in one dimension (such as eye height) do not always have the longest measurement for all other dimensions (for example, forward grip reach). What this means is that for practical purposes, each anthropometric variable could be considered independent of others. (Note that the level of correlation among different metrics can vary - for example, different attributes of the hand are closely correlated to each other, but measurements of different limbs are weakly associated.) So when we are setting an eye height that accommodates the lower 5% to upper 95% of that metric, and then a forward grip reach that accommodates the lower 5% to upper 95% for that particular metric, we are actually talking about two different groups of people. Only a subset of people who fall within the eye height range will also fall within the reach range, albeit a large subset, but below the 90% of the population we are striving to include.

Practical Solutions

There are some analytical methods for more effectively addressing these issues mathematically, but that's beyond the scope of discussion (for those interested, see *Guidelines for Using Anthropometric Data in Product Design*). In practical terms there are three solution approaches: design multiple sizes, adjustability and satisficing.

Multiple sizes, as it implies, creates a range of models, where each is targeted at a specific subset of the user population. The most extreme example of this (aside from bespoke, individualized designs) comes from clothing and footwear, where there are literally dozens of sizes and variations to enable a relatively close fit for the vast majority of the population. For products such as furniture, this may be limited to three or four sizes, better known as small, medium and large. In fact, this was Herman Miller's solution to the chair fit problem - creating three different sizes allowed for fit of 95% of the population between the smallest 1st percent and highest 99 percent - a greater range than they had originally intended. During the design of the airport kiosk that we discussed in part 1, one of the early proposed solutions was to create a two-sided kiosk with a "low" and "high" screen positions that could comfortably suit a wide range of users.

Adjustability is really a special case of multiple sizes where all possible sizes are provided by a single version of the product that the user (or an installer) modifies the fit at installation or during use. Most of us are familiar with adjusting the driver's seat in a car. These seats are not infinitely adjustable, but typically have three or more control points that can lead to a very wide range of positions, within the available space constraints. The downsides of adjustability are cost, reliability, and the extra work placed on the user to adjust the fit. Note, that many users may not always set the best fit for themselves.

Satisficing, is coming up with a single solution that fits the broadest range of users. In practice this tends to skew towards the smaller or shorter end of users because, larger users can always bend and smaller users may have physical limitations due to age or disability that take priority (legal and otherwise). Most designs for public spaces will take this approach, as in elevators, water fountains and ATMs.

Prototyping for Fit

Whether designing a single solution or multiple sizes, it is important to follow a user-centered design process. As in interaction design, prototyping can take many forms, depending on your goals and need for fidelity at each stage of the design process. For example, if the initial goal was simply to conduct a real-world test of key dimensions, then a simple sticker on a wall could serve as a "prototype" for display position. For more detailed issues, such as task-specific grips on a tool handle, foam mock-ups can be created and evaluated.

A typical user-centered design process for ergonomic fit would follow these steps, presented in an abbreviated form here:

1. Define relevant populations (e.g. age range, nationality, sex)
2. Define key dimensions or variable for fit consideration (e.g. height, reach, weight, etc)
3. Determine boundary measures for each anthropometric dimension from reference data, from lower 5th to upper 95th percentile (keeping in mind that some dimensions, such as head clearance in a doorway, may be one-sided)
4. Compare referenced dimensions with existing real-world products for reality check
5. Apply dimensions to create mock-ups for initial, informal ergonomic feedback with users
6. Refine design(s) to create foam or similar low-fidelity mock-ups for fit evaluation
7. Continue to refine as needed/budgeted

Qualitative Observations Issues in Field Research

While interaction designers will typically lack special training in ergonomic assessment methods, most will have some degree of familiarity, if not significant experience with user-centered methods including contextual observation (aka ethnographic field research) and usability testing. All of these methods share objective observation as a common data gathering method, and really only vary in the particular variables or characteristics that are the subject of study. And while anthropometric data is intrinsically quantitative, qualitative observational research can be applied to identify ergonomic issues. With these factors in mind, I've developed a basic set of ergonomic observational criteria to use as guidelines when evaluating design fit. The guidelines are inspired by Stephen Pheasant's cardinal rules of anthropometrics, extended to qualitative field research.

Pheasant advised focusing on *Reach*, *Clearance*, *Posture* and *Strength*. I'll explain how these can be applied to a consumer electronics device, the InterAction Labs SQWEEZE Game Controller. The SQWEEZE is an accessory to the Nintendo Wii - inserting a Wii controller into the SQWEEZE unit allows the user to apply push/pull forces for gaming - think of drawing a bow string to shoot an arrow, for example. While the SQWEEZE was well designed by ergonomics standards, it makes for a good example for explaining the four anthropometric characteristics:

- *Reach* typically refers to extending the arms and finger for effective control without over-extension. In the case of the airport Kiosk discussed in Parts 1 & 2 there's a clear potential for placing the touch screen at a height or distance that would be difficult for some people to access effectively. That type of reach is a non-issue for handheld devices like the SQWEEZE, but other types of reach can come into play. In the case of two-handed devices, the distance between the handles needs to be appropriately set to accommodate a comfortable grip. For the SQWEEZE, this distance actually varied between the push and pull positions as the handles flexed inward and outward respectively. Similarly, the diameter of the handles affects the user's ability to adequately wrap his or her fingers around them; a smaller-scale, but just as important, reach issue.
- While reach is about making sure things are not too far away, *clearance* is primarily focused on making sure things aren't too close together. In interaction

design terms, we might think of this as literal "white space". There needs to be adequate room for the hands to move around the handles without bumping into anything, constraining usability or performance.

- We tend to think of *posture* as a full-body issue; standing upright or bending. But in fact posture, defined as deviation from a natural, comfortable position, can be examined at the level of a specific limb or limb-segment. In handheld controllers, wrist posture is frequently the factor of interest. A design that forces the joints into contorted, unconformable positions, particularly for extended periods, is an ergonomic failure.
- *Strength* was particularly important for the SQWEEZE as it's essentially a force transfer device. Testing with children indicated the device should not exceed 2.5lbs, but it also had to withstand up to 150lbs of crushing and pulling - the strength of a 90th percentile male. In more general terms, designs should avoid requiring significant exertion by the user, but need to have sufficient resistance to provide feedback and avoid accidental triggering, for example as on a mobile phone keypad.



I've just scratched the surface of these four key ergonomic factors, but I want to re-enforce a couple of critical issues to keep in mind. First, when we talk about a particular factor, it's important to consider it at multiple levels of scale. In the case of posture, we might look broadly at how someone approaches a kiosk from an overall body perspective, but then focus more narrowly on the deviation of the hands and fingers. Second, these factors are not independent of each other - in fact they are highly co-influential. For example, if there is limited visual access, then a user may change his or her body and limb postures to accommodate improved field-of-view, but in doing so, increase the extent of reach and reduce the effective transfer strength.

Last, but not least, I add a fifth factor which goes beyond the physical, to the perceptual and cognitive: *Feedback*. Feedback refers to the user's ability to receive input on the impact of their actions on the interface or system. For the SQWEEZE this can mean the tactile, visual and even audible mechanical feedback that corresponds with using the device. For a touch screen kiosk, there is the perceived resistance of the touch service, and the feedback from the software responses.

Putting all this together, a person conducting observational research can use these five factors as a checklist for identifying potential ergonomic problems in real-time, or post-hoc (e.g. with video review).

As a mnemonic aid, putting Feedback together with the other four ergonomic factors (Reach, Clearance, Posture and Strength), gives us FRCPS, or FoRCePS. This was actually created as a mental cue during surgical observations, thus the clinical abbreviations. I'm certainly open to more approachable re-combinations of the letters.

Measured vs. Perceived Fit

In more formal assessment situations, such as usability testing, there are a number of quantitative methods for measuring fit and identifying ergonomic problems or risks. But what seems well-designed on paper doesn't always result in well-received or usable. I've observed numerous situations where the "technical" ergonomic requirements of a design would suggest a good fit, but in reality, the majority of users preferred an alternative. There are various reasons for this ranging from individual differences, to preference for the familiar, to the influence of aesthetic design. It's not the reason for these outcomes that matters so much as the need to capture this input. In other words, it's just as important to measure subjective or perceived fit and comfort, as it is to measure anthropometric fidelity.

Recently, a number of surveys and guidelines have become available for measuring perceived comfort (I realized *perceived* comfort is redundant, but I'm including it for clarity). For example, Kuijt-Evers, Vink & De Looze² present a basic survey for hand tool comfort that covers factors from ease of use, to performance to...blisters. In practice, it's helpful to use a vetted survey like this as a starting point, and then add and subtract questions based on the particular needs of your product, users and tasks, paying attention to the FoRCePS issues described above. As with any user-research study, piloting and iterating the usability testing approach is as important as iterating the design itself.

² <http://tinyurl.com/nft2ma>