

Quantifying Degree of Goal Directedness in Document Navigation: Application to the Evaluation of the Perspective-Drag Technique

Yves Guiard¹Yangzhou Du²Olivier Chapuis²

¹ LMP – Univ. de la Mediterranee & CNRS
F-13288 Marseille, France
yves.guiard@univmed.fr

² LRI – Univ. Paris-Sud & CNRS, INRIA
Bât. 490, F-91405 Orsay, France
{du, chapuis}@lri.fr

ABSTRACT

This article pursues a two-fold goal. First we introduce *degree of goal directedness (DGD)*, a novel quantitative dimension for the taxonomy of navigation tasks in general. As an attempt to operationalize the *DGD* concept in the context of electronic documents navigation, we introduce the *serial target-acquisition (STA)* experimental paradigm. We suggest that *DGD* and the *STA* paradigm may usefully enrich the conceptual toolkit of HCI research for the evaluation of navigation techniques. Our second goal is to illustrate the utility of the *DGD* concept by showing with a concrete example, *Perspective Drag*, the refinement it allows in evaluating navigation techniques. We report data obtained from two experiments with the *STA* paradigm that cast light on what *Perspective Drag* is specifically good for: it is particularly suitable in realistic task contexts where navigation is less than 100% directed by its terminal goal, that is, where the user wants not only to reach a particular item but also to pick up information from the document during document traversal.

Author Keywords

Document navigation, Interactive multi-scale visualization, Goal-directed behavior, Fitts' law, Task taxonomies.

ACM Classification Keywords

H.5.2. [User Interfaces]: Interaction styles; I.3.6. [Methodology and Techniques]: Interaction techniques.

INTRODUCTION

One may think of a broad diversity of possible reasons why users navigate a document in a graphical user interface (GUI). They may (i) browse the document aimlessly, just hoping to discover something interesting, or (ii) search the document for a specific type of contents, say illustrations,

filtering out everything else, or (iii) want to get to some well identified and well localized item, say they want to reach a certain reference line in the final page of the document ten pages down, to do some editing.

What these examples have in common is that users need to move their view across document space, and so all three cases qualify for the document navigation metaphor. But to classify them, we need more than one taxonomic dimension. In the three cases users obviously look for information while traversing document space, but the three examples differ along at least two dimensions: whether or not the users know the nature of the information of interest and whether or not they know the location of this information in document space.

Our option in this research on document navigation is to leave aside variables that relate to the nature of information and to focus on those regarding information *location*. For simplicity we assume that all the items the user is looking for are of the same sort and of the same degree of interest. This simplification takes us away from such approaches as Furnas' [5,6], whose a priori importance functions encompass both of these aspects, as well as Pirolli and Card's [16], who proposed an optimal information foraging model where the information value of individual items is a critical variable to be evaluated against their discovery and handling costs.

Speaking taxonomy, what task variables are we then left with? One variable is the extent to which the user is aware of the location of the items of interest. Again for simplicity, we will fix this variable—we shall assume the user knows where the interesting information rests in document space. So, having confined our inquiry in the narrow case of users who know *what* they are looking for and *where* interesting things are (i.e., the document to navigate is familiar), we end up with one task variable, the *layout*, in document-space, of the information the user wants to visit. Although the information layout is still a multi-variable concept, below it will become apparent that it is neat enough conceptually for use in quantitative modeling and experimentation.

Although we are interested in electronic document navigation, our emphasis in this paper is more on the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2007, April 28–May 3, 2007, San Jose, California, USA.
Copyright 2007 ACM 978-1-59593-593-9/07/0004...\$5.00.

structure of documents (the environment for navigation) than the dynamics of the navigation process. Degree of goal directedness in our view is a task characteristic, and tasks can be analyzed and classified independently of the way in which they are carried out.

Many studies have tried to understand information acquisition in large hierarchical databases, e.g. [11,16]. Here we will consider the simpler case in which the information of interest is contained in a single open document. We will actually take the navigation metaphor fairly literally: navigating a document here means traveling one's virtual camera in Euclidian 3D space, hovering over the flat landscape of a document, in such a way as to shift spatially and rescale one's view of the document.

DEGREE OF GOAL DIRECTEDNESS AND FITTS' TARGET-ACQUISITION PARADIGM

Fitts' target-acquisition paradigm [3] has long been recognized in HCI to be a useful and robust tool for rigorously evaluating both simple [13] and multi-scale [8] navigation techniques. However, to correctly appreciate the potentialities of the paradigm with regard to the evaluation of navigation techniques it should be kept in mind that, as strongly emphasized in [15], this paradigm captures a *limiting case*.

Fitts' Paradigm and What It Fails to Capture

The concept of *goal-directed behavior* is a classic of scientific psychology. It has a long history in ethology, which for example has endeavored to understand the homing behavior of pigeons and other non-human animals [19]. Turning to human experimental psychology, we have the generic concept of aimed movement, whose study was provided half a century ago by Paul Fitts with a stable experimental paradigm, the so-called pointing, or target-acquisition paradigm [3].

The target-acquisition paradigm was initially introduced in HCI to study simple pointing gestures within a view implicitly assumed to be stationary [2,13,18]. It is useful to recall that in recent years the paradigm has been extended to the case in which the user wants to visualize some remote, currently not visible target object, by moving the view, a task category which we have termed *view pointing* [8]. Another noteworthy development that has recently taken place in HCI about Fitts' task was the demonstration that one can still use Fitts' paradigm, and demonstrate Fitts' law, in interfaces that allow the user to freely rescale the view (e.g., with a zoom): Fitts' paradigm successfully accommodates the case of *multi-scale* target acquisition. This is the case we will be considering exclusively here.

The limits of Fitts' paradigm derive from its very strength. This paradigm focuses accurately on a very specific kind of behavior, and being accurate has a cost. Fitts' law experiments typically ask participants to traverse empty spaces to reach some specified target object, but we must

recall that real documents, be they physical or electronic, are covered with information.

Degree of Goal Directedness: Principle of Quantification

With his concepts of *a priori importance* (API) and *degree of interest*, Furnas [5,6] has made it clear that document navigation cannot be understood without consideration of the *distribution*, in document space, of the information the user is interested in. The idea we introduce below is a limited response to this suggestion. To start with, notice that, however one defines information (e.g., using Shannon's measurement), that quantity must be supposed to accumulate during the navigation process. Along lines reminiscent of Pirolli and Card [16], our idea of goal directedness in navigation will be defined based on the rate at which cumulated information payoff rises as the view progresses en route to some terminal target.

In the case of goal-directed tasks as scientific psychology traditionally defines them [3,8,13,19], we get an all-or-none relation, as shown in Fig. 1 (thick, horizontally-inverted L curve): cumulated payoff remains zero during the whole travel and jumps to 1 as the target is reached. It is the near boundary of the target, at the end of the path—conceptually an extensionless point—that delivers the whole payoff. In this case, quite clearly exemplified by a Fitts task, the movement can be said to be 100% directed by its terminal goal.




Fig. 1. Principle of quantification of degree of goal-directedness in document navigation over some definite path.

Another special category of navigation tasks is that corresponding in Fig. 1 to the dotted line running along the diagonal. Here every single point along the space portion traversed by the view is of interest and so cumulated payoff rises at a constant rate: covering $x\%$ of the distance means obtaining $x\%$ of the total payoff expected from the navigation act. The terminal point of the traversed space does not have any special status, its influence being totally diluted. Let us label this special case as pure (totally aimless) exploration.

These two cases differ along the dimension we call the *degree of goal-directedness (DGD)* of the navigation task. Let us now define for any curve a quantitative expression of *DGD*, using some dimensionless ratio which we want to vary from 0 to 1. Consider the surface area (A_C) under some arbitrary monotone curve C . Since the curve is contained in a square of side 1, A_C varies in the $[0,1]$ interval. Obviously, if the payoff is to be delivered as a whole in the end, as in a Fitts task, we have $A_C=0$; symmetrically if the whole payoff is to be delivered at the start we get $A_C=1$, but that case, which means navigation abortion, should be ignored. To obtain a valid measure of *DGD*, we must scale A_C to some reference value that makes sense in our modeling of navigation tasks. We shall use as our reference the surface area under the diagonal, equal to $\frac{1}{2}$, because the diagonal corresponds precisely to the case in which *DGD* is indeed minimal. We can now define the degree of goal directedness for any curve C as

$$DGD_C = 1 - (A_C / \frac{1}{2}) = 1 - 2 A_C. \quad (1)$$

As desired, we obtain $DGD_C = 1$ or 100% for a totally goal-directed navigation task and $DGD_C = 0$ or 0% for a totally aimless task.

There must exist a broad variety of aimed tasks that score between 0% and 100% on the degree of goal-directedness continuum. Fig. 1 provides an arbitrarily chosen example (gray curve), in which we see that the rate of change of the payoff with respect to the distance covered increases monotonically as the view progresses towards the terminal target. Such a case can be characterized by saying that the influence of the target is partially diluted—it is only *to some extent* that the navigation is goal-directed. The reason why this intermediate case, whose *DGD* falls somewhere within the $[0,1]$ interval, is worthy of consideration in the context of HCI is because it is liable to capture a large proportion of the document navigation tasks that computer users actually carry out everyday. Even though the 100% goal-directed mode of navigation is remarkably suitable for experimentation, as demonstrated by the large body of solid knowledge obtained in HCI research thanks to Fitts' methodology [2,8,18], it is hard to believe that when computer users navigate some document to get to some specific spot, be it perfectly well defined, the interest of the contents they encounter during document traversal is *zero*. HCI reality, we shall argue, is more faithfully modeled with the intermediate curve rather than the inverted-L curve of Fig. 1.

The foregoing was aimed to suggest that the classic task typology that distinguishes, on a nominal scale of measurement, such familiar entities as exploration, search, and goal reaching is amenable to a quantitative treatment. It is obvious from Fig. 1 that, accepting a few assumptions, the variable of goal-directedness can be raised to the level of a ratio scale. Here is one important preliminary assumption we will be making: we will only consider navigation tasks that are goal directed to some extent—that is, tasks for which a *terminal* goal is identified. This

excludes from the scope of this modeling, for example, the case of a navigation task in which, although there is a distant spot to reach, most of the payoff is to be delivered after a small amount of navigation because the majority of the interesting stuff is located just after the start. In such a case we believe the terminal spot should be regarded as an outlier, rather than the terminal goal of the navigation. Using the logic of Fig. 1, this means we want to ignore the case of concave-down curves. We will consider only *concave-up* curves, up to the diagonal line.

The reader will perhaps notice that the curves of Fig. 1 are reminiscent of the Lorentz curve of econometrists [17]. A Lorentz curve obtains when the cumulated percentage of wealth owned by a given population is plotted as a function of the cumulated percentage of households in that population, households being ordered from the poorest to the richest. While the extreme curves are the diagonal (perfect equality) and the inverted L (total inequality) the Lorentz curve technique makes it possible to express graphically (and quantitatively, using the Gini index [17]), the characteristic degree of inequality of some population by the degree of concave-up curvature of the obtained curve. Such an analogy makes sense. What we want to characterize in the present study of navigation is the inequality of the distribution of interest [6] in a portion of document space that one wants to navigate, en route to some target. Perfect equality and inequality are intellectual fictions in econometric modeling; likewise, for real HCI tasks the limiting cases of totally aimless exploration and 100% goal-directed navigation, notwithstanding their conceptual utility, are theoretical limiting cases.

A First Attempt at Making *DGD* Operational: The Serial Target-Acquisition Paradigm

This section aims to define one operationally accurate, and hence experimentally usable implementation of the *DGD* concept. Fig. 2 shows three possible arrangements of a number of target objects that have been placed in some document displayed as a linear array of pages—so reasoning here is 1D. Assume the targets are lines of text, shown as gray patches in the figure, and they need to be visualized, that is, captured one after the other with the view (the horizontal rectangle shown over the bottom page from where navigation is supposed to start).

Curve A of Fig. 2 corresponds to usual view pointing [8]: the task involves a single target, and so for the participant the cumulated payoff of navigating the document will remain zero all along the distance covered and suddenly jump to 100% as the target is reached. The surface area under the curve is zero and hence, by Equation 1, *DGD* is 100%. Curve B corresponds to the case of there being several targets (here $N = 3$). Cumulated payoff switches to $1/3$, $2/3$, and 100% as the first, second, and final target is reached. Measuring the surface area under the curve, one finds a *DGD* of 75%. Finally, we get Curve C when every single line of the document calls for a visit—i.e., interest is distributed homogeneously all along the travel ($DGD = 0\%$).

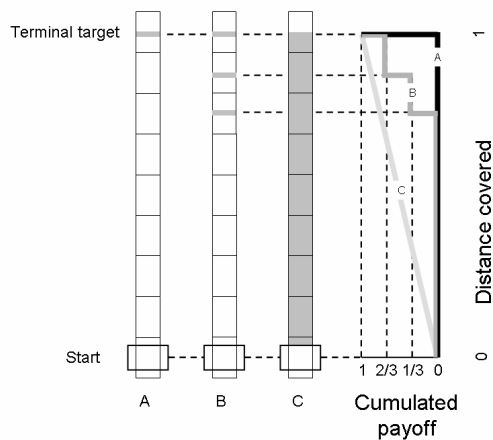


Fig. 2. The serial target-acquisition paradigm.

The words “target acquisition” here essentially refer to view navigation through space and scale [6, 8] rather than simple (single-scale) pointing. Of interest in this study is the navigation of documents that are large enough to raise the focus + context problem and therefore to demand multi-scale interaction. Target distance D and target width W can be defined just as accurately as they are in traditional Fitts law experiments [2, 3, 13], and therefore it is still possible to compute an index of difficulty $ID = \log_2(D/W + 1)$ for the navigation. However, movement time (MT), the main dependent variable of a Fitts’ law experiment, corresponds here to the time it takes people to visit, that is, to adequately *visualize* the target (which, to reiterate, typically requires both space and scale manipulations), regardless of whether or not they must eventually select (e.g., click) the target.

The value of the STA paradigm, in our view, rests in the fact that it captures some important characteristics of real document navigation tasks while being defined in a rigorous and concrete way. It is easy to find examples of real HCI tasks where the user needs to systematically visit a whole series of spots in a document: just think of a user who needs to check the formatting of the legends of all the figures (or tables or headings) of a document (s)he is editing. This means performing a sequence of navigation steps up to the final item.

Indeed, a proportion of these tasks can be handled efficiently using keyboard shortcuts or automatic search mechanism (e.g., skipping from reference to reference by searching for the “[” character) but fortunately interface design also provides users, along with automatic search mechanisms, with efficient techniques to travel their documents *actively*.

APPLYING THE DGD MEASURE: A COMPARATIVE EVALUATION OF PERSPECTIVE DRAG

We now turn to the second facet of this contribution. What follows is *not* meant to provide empirical support for the DGD concept. Rather, we aim to suggest, with one example illustration, that the DGD concept, together with its metric,

is helpful for thinking of navigation tasks and evaluating interfaces. We will offer an illustration of the heuristic value for HCI research of the *DGD* concept, using a concrete example, namely the evaluation of the Perspective-Drag (PD) document-navigation technique. By the same token, we will try to convince the reader that, based on the evidence we collected with the *DGD* approach, it would be quite useful to provide users with that new tool for the navigation of many categories of documents.

The Perspective-Drag Technique¹

In the state of the art technology, the camera which stands at the core of the GUI has only three degrees of freedom (DOF), all translations [4]. The user can translate the camera in 3D space relative to the flat document to obtain zooming (manipulating observation *altitude*), and scrolling (manipulating camera *longitude* and *latitude*). Suppose you provide the user with control over one more DOF of the virtual camera, a rotation DOF (pitch). Now users are able, if and when they will, to obtain a *perspective* view of their document by tilting their camera, as shown in Fig. 3.



Fig. 3. An example of the view that users can obtain over a document with the PD technique (the screen cursor, not shown here, is the hand symbol of Acrobat Reader).

A number of converging analytic arguments point to the usefulness of a camera-tilt facility in GUIs. Here are a few that we find very compelling. Geometrically speaking, the perspective projection, which involves a progressive, highly nonlinear, and quite familiar variation of visualization scale across the view, offers a powerful and elegant solution to the focus + context problem that constitutes the central concern of information visualization research [6,14].

Tilting one’s viewpoint in the direction of interest is a more natural and efficient way of *looking for* the location one wants to reach in document space, prior to navigation departure, than translating one’s viewpoint. In particular,

¹ An interactive hands-on prototype of the Perspective-Drag technique is publicly available at <http://www.laps.univ-mrs.fr/~guiard/pvi/>.

compared with zooming-out, which often forces irrelevant regions of the document as well as empty background space to enter the view, a camera tilt optimizes the utilization of screen pixels [10].

During the course of a pan or scroll, a tilted view over the document is a necessary condition for the user to obtain *prospective* visual information. As demonstrated by psychologist James J. Gibson [7] through an analysis of the optic flow fields that are induced by motion of an observation point, one needs to *look ahead*, rather than at right angle to the navigated surface, to obtain the anticipatory kind of visual information one needs to steer one's progression [10].

It should be emphasized that scrolling with the familiar drag gesture (e.g., of Acrobat Reader) becomes highly *nonlinear* when applied to a document viewed in perspective—and this nonlinearity is an asset for navigation control. Assuming the camera is tilted to the beginning of the document as in Fig. 3, suppose you want to see the detail of some remote region (appearing in the upper part of the view). Grasp that region with your screen cursor, and drag it to you, downward: as your grasping cursor (Acrobat's little hand) moves downward in the view the display-control (DC) gain for document motion will drop exponentially. This means that even with a constant mouse velocity the approaching speed of the grasped region, extremely high at first, will eventually die out as the target gets right at the desired location down the view, right at the appropriate scale [10]. Of course, if you grasp a near spot in the document and push it away (obtaining a 'Star-Wars' sort of scroll) the DC gain will *increase* nonlinearly, so this must be done cautiously in small steps. Finally, note that if you have tilted your camera the other way round—to see the bottom of the document from the top—the display will look less familiar (the document surface appearing like a ceiling rather than a floor), but all the above properties will hold by up-down symmetry.

Recent Evaluation Work on Perspective Drag

In a recent study [10] we submitted the PD technique to a formal evaluation using Fitts' target acquisition task, in comparison with what we designate as the *Zoom-and-Drag* (Z&D) technique of widely used applications like Acrobat Reader. The results (Fig. 4) revealed that, strictly speaking, movement time (*MT*) does not obey Fitts' law with PD: that is, expressing *MT* as a function of the index of difficulty $ID = \log_2(D/W+1)$, all our individual curves for this technique exhibited some degree of concave-up curvature. Beside this nonlinearity, which indeed can be predicted mathematically as an effect of the nonlinearity of perspective projection [10], the most clear-cut finding was the intercept and slope difference, leading to a cross-over interaction pattern: specifically, the Z&D and the PD curves crossed at about an *ID* of 12 bits, suggesting that while the PD technique was weaker than its competitor for $ID > 12$ bits, it was stronger for $ID < 12$ bits—indeed, for $ID = 9$ bits

PD already outperformed Z&D significantly. Now it should be realized that for goal-directed navigation an *ID* of over 12 bits, which corresponds to a *D/W* ratio of over 4,000 (e.g., reaching a target line at a distance of 47 pages, using the SIGCHI format), is unlikely to be met outside Fitts' law research laboratories. Thus, as far as multi-scale pointing in real situations is concerned, it is safe to conclude from the results of Fig. 4 that PD is in general more efficient than Z&D.




Fig. 4. Results of our previous evaluation [10] of the PD and the Z&D techniques, using a 100% goal-directed task.

EXPERIMENT 1 : PERSPECTIVE-DRAG VS. ZOOM-AND-DRAG IN SERIAL TARGET ACQUISITION

Experiment 1 simply aimed to test our expectation, deduced from our previous tests of PD technique with the conventional, 100% goal-directed Fitts task, that for a more realistic task requiring not simply to reach a remotely located target, but also to visit a number of other spots during the traversal, PD becomes more efficient than the Z&D technique available in many of our familiar applications. We wanted to show that with a task whose *DGD* is substantially less than 100% the PD technique does enhance navigation performance, a fact worth establishing if it is recognized that a tool which reveals good for a task that one has to do frequently is worth having at one's disposal.

Method of Experiment 1

The document our participants were asked to navigate was obtained from a richly illustrated five-page article from *Science* magazine.² We concatenated 12 copies of this article to build a 60-page long document, which our experimental interface displayed as a vertical linear array, as in Acrobat Reader (full-screen mode, effective display resolution = 1,270 x 909 pixels).

We used the above-described STA task. Participants were to check as fast as possible whether spelling was correct ("next") or wrong ("nest") for each of the 12 occurrences of

² *Science* 309, 23 Sept. 2005, 1989-93.

a specified word in a specified line. The periodical line of interest was easy to find, being located just above a conspicuously colored figure (periodical too, like everything else in this document) that was recognizable from quite a distance. Participants were to proceed from the bottom to the top of the document (i.e., from the 12th occurrence of the target line upward). We had edited the document so that the probability of encountering a wrongly spelt word was $\frac{1}{2}$, and so each instance of the word had to be carefully checked. Depending on whether the target word was correct or wrong, the participant was to produce a double-click with the left or right button of the mouse. One document required a total of 12 double clicks, yielding 11 time measurements.

The task involved reaching targets and had a definite terminal goal (the 12th occurrence of the problematic line), thus qualifying as an aimed task. However, it differed markedly from conventional Fitts tasks in that its *DGD* was low, amounting to a modest 9%. Note also that this was not, strictly speaking, a pointing task. The spatial tolerance was extremely large for the double click, which had to be placed anywhere within the currently checked page. So two sorts of errors were possible: classification errors (double clicking the wrong mouse button, meaning an error in judging word spelling) and navigation errors (double-clicking a page different from that which had to be checked, meaning a failure to follow the sequence from bottom to top of the document). As an incentive for cautiousness, any error caused the program to restart the trial from the bottom of the document.

At a pinch this task could have been done by just scrolling the document, at some well chosen constant level of scale, but that would have required an uncomfortable amount of mouse scroll. Considering two consecutive target words, the ratio of target distance and target size was actually large enough ($D=1,126$ px and $W=16$ px, yielding a Fitts' *ID* of 6.2 bits) that multi-scale visualization was quite helpful.

Participants were offered two multi-scale techniques. The **Z&D** technique we used was an emulation of the well-known Acrobat Reader technique, where document viewing is fixedly perpendicular. To check each critical line of our 60-page document, starting from its end, with Z&D the participant had to (1) zoom-out with the mouse wheel until the target line entered the view, (2) mouse drag the document to bring that line close to view center, i.e., the zoom expansion focus, (3) zoom-in with the wheel to magnify the words up to reading size, and finally (4) double-click to indicate whether word spelling was correct or wrong. Thus the Z&D technique involved the coordinated manipulation of camera altitude (zoom) and latitude (forward scroll) and required $12 \times 4 = 48$ elemental operations per document.

With the **PD** technique, whose crucial feature is the camera tilting facility, the participants first prepared their navigation by setting their viewing angle (camera tilt) and

observation altitude (zoom) so as to obtain a perspective view of the whole document from the bottom (as in Fig. 3). Then the task was carried out by means of two elemental operations: (1) mouse drag the next target line downward to oneself, thus magnifying the word up to reading size (typically a single drag gesture sufficed to produce the 5-page scroll) and (2) double click to indicate one's decision about word spelling. Thus the PD technique required $12 \times 2 = 24$ elemental operations per document.

Sixteen unpaid volunteers participated in two short sessions (about 15mn each) each consisting of three complete trials with one of the two techniques (producing a total of $3 \times 11 = 33$ navigation acts per participant). Even-numbered participants used the Z&D technique for the first session and the PD technique for the second, the order being reversed for odd-numbered participants.

Results of Experiment 1

Error rates

On average over all 16 participants, the percentage of classification errors was 0.9% and 1.0% for the Z&D and the PD conditions, respectively. The corresponding figures for navigation errors were 1.4% and 0.9%. Neither difference was significant.

Navigation Speed

Since the experiment just consisted of comparing two navigation techniques, let us simply consider how movement time (*MT*), defined as the time elapsed between two consecutive double clicks, evolved with short-term practice over the duration of the experiment for the two conditions.

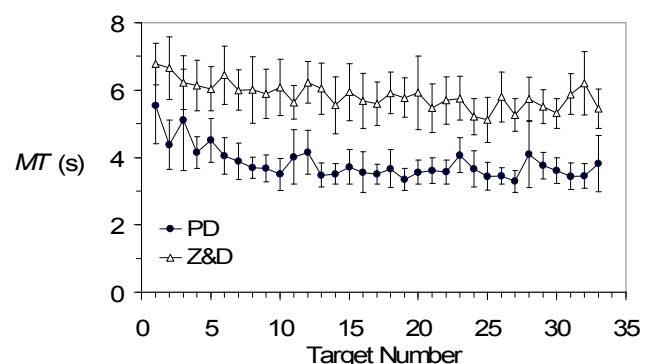


Fig. 5. Performance with the PD and Z&D techniques over all the movements recorded in Experiment 1.

Fig. 5 concatenates the three successive trials to show the whole sequence of 33 elemental *MT*s, each averaged over the 16 participants (in this figure and others below error bars represent 95% confidence intervals based on between-participant standard deviations). It is quite clear that PD surpassed the usual Z&D technique quite consistently. On average over the last two trials (from the 12th to the 33rd move), where performances were roughly stabilized, the

time saving allowed by PD relative to Z&D was 36.2%. All 16 participants exhibited an advantage of PD over Z&D ($p < .001$, 1-tailed sign test). Notice that although the PD technique was entirely new to our participants, a few double clicks sufficed for them to command it and fully exploit its benefits.

Subjective Evaluation

Unsurprisingly given the performance data, all 16 participants reported they preferred PD over Z&D for that task ($p < .001$).

Discussion

Thus the data of Experiment 1 quite clearly confirmed our expectation that the PD technique would do far better than the Z&D technique which we use almost every day in a task where, unlike Fitts' task, the items of interest for navigation were evenly spread along the path, yielding a *DGD* far below 100%.

The explanation of the superiority of PD for this sort of task is straightforward. To examine a series of detailed items scattered in a large document users of the Z&D technique must have recourse, in a carefully coordinated way, to both zooming and scrolling [1,9]. PD navigation, in contrast, requires just scrolling. Once the altitude and tilt of the camera have been appropriately set, a sequence of mouse drags suffices for a finely-grained inspection of the whole document, however large. Thus allowing users to tilt their virtual camera to obtain a convenient perspective view results in a simplification of navigation, which now involves a single input DOF.

EXPERIMENT 2: INFORMATION DENSITY, INFORMATION DISTRIBUTION, AND REACHING DIFFICULTY

Using the same task (checking, under speeding instructions, the spelling of several occurrences of a target word) and the same two techniques as in Experiment 1, we proceeded to manipulate the *DGD* variable directly with a quantitative approach. We predicted an interaction: the performance superiority of PD over the Z&D technique should increase as the *DGD* is reduced, that is, as the interesting information gradually spreads itself in the document.

One can imagine several possible ways of varying *DGD* in a navigation task. We chose two different manipulations. One consisted of varying *information distribution* in the document, using a constant number ($N=6$) of target items but different spatial arrangements, as described in Table 1A. For this manipulation we designed three documents: in Doc#1 the six copies of the target page were evenly distributed, one every 6th page, yielding a low *DGD* of 20%. In Doc#2 the target pages tended to squeeze toward the beginning of the documents (actually the ending of the navigation, from the participant's viewpoint, since navigation was to proceed upward), with an arithmetic reduction of target spacing ($DGD=47%$). In Doc#3, the squeeze was more pronounced, with a geometric reduction of the spacing ($DGD=65%$).

A					
Doc.	N	Layout	ID _c	DGD	Item location (page #)
#1	6	Equidist.	43.63	0.20	{31, 25, 19, 13, 7, 1}
#2	6	Arithm.	42.62	0.47	{31, 21, 13, 7, 3, 1}
#3	6	Geom.	40.63	0.65	{31, 16, 8, 4, 2, 1}

B					
Doc.	N	Layout	ID _c	DGD	Item location (page #)
#4	3	Equidist.	20.09	0.50	{31, 16, 1}
#1	6	Equidist.	43.63	0.20	{31, 25, 19, 13, 7, 1}
#5	11	Equidist.	77.29	0.10	{31, 28, 25, 22, 19, 16, 13, 10, 7, 4, 1}

Table 1. Varying information distribution (A) and information density (B) in the document.

The other manipulation consisted of varying *information density* in the document, by varying the number of target items, keeping them equidistant, as explained in Table 1B. We again used three documents, but we reused Doc#1, with its six copies of the target page, to be compared with Doc#4, which had fewer target pages ($N=3$, $DGD=50%$) and Doc#5, which had more ($N=11$, $DGD=10%$).

An Intricate Relation Between *DGD* and Fitts' *ID*

Notice, however, that these two manipulations can be described in another way, adopting a Fitts' law approach. We used a serial target acquisition task and so total navigation time can be predicted as the sum of a sequence of elemental pointing times, each presumably determined by its *ID*, according to Fitts' law. Now the sum of all elemental *ID*s (the cumulated *ID*, or *ID_c*) varies in a non-trivial manner with target number and distribution (Fig. 6).




Fig. 6. Co-variation of *DGD* and Fitts' *ID* in the manipulation of information density and distribution of Exp. 2.

Our manipulation of the distribution of a fixed number of targets entails little or no change in the *ID_c*, meaning that the effect of the *DGD* factor can be experimentally isolated with this method. However, our manipulation of the number of evenly distributed targets (i.e., *information density*) involves a factor confound: the higher the *DGD* (i.e., the fewer target items in the document), the lower the *ID_c*. Thus, we found it safe in Experiment 2 to have recourse to both methods of manipulating *DGD*.

Attention should be called to the necessity of carefully distinguishing, in a STA task, the *ID_c* from mean elemental difficulty: if a document offers more intermediate targets,

the task's ID_c will increase, but obviously the mean difficulty of elemental reaching acts will decrease.

Method of Experiment 2

The Document we used in this experiment had 32 pages taken from the same *Science* magazine article as in Experiment 1. The pages were re-ordered in such ways that the copy of the target page (that containing the word with a possible spelling error) would appear in the sequence either 3, 6, or 11 times (the recurrence being strictly periodic spatially) or with a variable distribution: equidistance, arithmetic or geometric squeeze (in this case with a constant number of occurrences, $N = 5$).

Note that prior to starting to navigate each document participants were shown graphically where the target pages were to be found in the document. So their task was definitely not a search. Rather, like users who navigate some familiar document, they had to reach a series of items whose number and locations were known to them.

Twelve unpaid volunteers participated in a single session (ca 30mn) the first and second half of which were assigned to the Z&D and the PD techniques, the order being reversed for even- and odd-numbered participants. For each technique they were asked to check each of the five documents five times, following a 5 x 5 Latin square. This experiment involved a total of 135 word checks per participant and per technique (135 x 2 techniques x 12 participants = 3,240 checks overall).

Results and Discussion of Experiment 2

Error rate

Unsurprisingly, keeping in mind our error-penalization method, errors were again rare in Experiment 2. While for word classification the error rate was 0.7% with both techniques, for document navigation the error rate was 1.0% with Z&D and 0.5% for PD, a non-significant difference.

Navigation Speed

Document navigation time, computed from the first to the last double click, was strongly dependent on the technique. On average over all document types, the data showed a time saving of 31.3% (17.93s relative to 26.08s) with PD relative to Z&D.

Our main concern here is the impact of our two methods of manipulating DGD . As shown in Fig. 7, our manipulation of *information distribution* was ineffective. We obtained no evidence, for either technique, that for a constant number of items, the navigation time was dependent on the layout of target items—despite the substantial amount of DGD variation entailed by this manipulation (see Table 1A).




Fig. 7. Effect of information distribution on document navigation time for each technique.




Fig. 8. Effect of information density on document navigation time, for each technique.




Fig. 9. An alternative description of the results of Fig. 8.

In contrast, with both techniques we obtained a strong and consistent nonlinear effect of *information density*: as should be expected, the lower the DGD , the longer the time needed to traverse the document to the final item. The really informative finding is the strong interaction, visible in Fig. 8, between information density and technique, the slope being much steeper with Z&D than PD. So the suggestion is well in keeping with our expectation: the denser the relevant information in a document one has to navigate, the more pronounced the relative advantage of PD over the Z&D technique.

However, the possibility of an alternative description of the same data, due to the above-mentioned factor confound, should not be overlooked. As revealed in Fig. 9, it is also the case that document navigation time was linearly dependent on the ID_c , the sum of all elemental ID s in document traversal. The practical suggestion is the same—PD is more advisable as information density increases—but this is a different interpretation, which relies on an analysis of elemental target acquisition acts.

IMPLICATIONS FOR HCI DESIGN AND RESEARCH

One of our two main intentions in reporting this piece of research about the comparative merits of two techniques, one of which is new and the other familiar to every computer user, was to make a point about the method of evaluating navigation technique in HCI. It is now widely recognized in HCI research that there is no such thing as the intrinsic value of a navigation technique—indeed, we need multiple comparisons to judge techniques. However, it should be realized that comparativeness, in the above sense, is insufficient. Were an evaluation study performed on *all* the currently available multi-scale navigation techniques, it would provide no firm ground for a conclusion if it relied on a single task (say, reaching a very difficult target). The competing techniques might have ranked quite differently in a different task context [12]. Statistically speaking, this means that basic facts in an HCI evaluation study should have the form of *interaction patterns* involving several techniques and several tasks.

Even though comparisons of multi-scale navigation techniques based on Fitts' task do deliver the desirable form of results (as these comparisons produce one Fitts' law curve per technique), the Fitts' law approach is not immune to the criticism that target acquisition, however rigorous the definition of this behavioral category, is just one sort of task among many others [15]. We feel that obtaining a better sense of the whole spectrum of user motivations and contexts for navigating documents in homes and offices should be an important objective for future research in HCI. Hence our attempt, in the present study, to spell out degree of goal directedness, a dimension which we think might usefully enrich the current taxonomy of HCI tasks.

It will be interesting in the future to inquire into other ways of operationalizing the DGD concept experimentally. The STA paradigm which we chose in this study is a fairly straightforward implementation of the DGD concept, as is obvious when Fig. 2 is compared with Fig. 1—in both cases we have an information payoff that accumulates during progression to the final item, the increase being continuous in the abstract generality of Fig. 1 and discrete in the concrete instantiation of Fig. 2. The main advantage of the STA paradigm is that it makes it possible to leverage our knowledge about Fitts' law. However, there should be many other ways to implement the general idea that navigation acts, whether in the context of electronic documents or in the real world, are directed by their goal to variable extents.

An important issue that arose in this study is the applicability of Fitts' law analysis in a study aimed at quantitatively tackling goal directedness in document navigation. One might be tempted to conclude that Experiment 2 was essentially a failure, arguing that it failed to isolate the effect of our new independent variable, DGD , and that our data can *also* be understood, quite parsimoniously, in Fitts' law terms. Perhaps a more subtle lesson can be learned from this piece of research, recalling that DGD is a tentative new taxonomic dimension, a new way of characterizing navigation tasks—not a theory (in want of refutation or corroboration). Even though we have shown that ID_c and DGD co-vary more or less inextricably with each other in the situations we have examined, it should not be taken for granted that ultimately only one of these two factors should be the *real* cause of the effects we measured on navigation performance. The ID and the DGD concepts, which emerged from quite different approaches, are constructs that can peacefully coexist, each with its own heuristic potential. Many phenomena can be tackled from more than one theoretical standpoint (e.g., light as photons or waves), each having its own merits. For example, it is fair to recognize that the informative interaction pattern of Experiment 2 was anticipated from reflections about DGD , not Fitts' law. It is also the DGD idea that must be credited for the design of ID_c , an interesting new measure of navigation difficulty.

Nonetheless this experiment did produce further evidence that Fitts' target-acquisition paradigm is robust and broadly scoped. After our simple modification of Fitts' classic pointing paradigm, from which we were able to derive ID_c , to our knowledge a novel global measure of navigation difficulty, we found that Fitts' paradigm successfully accommodates still another case not foreseen by Fitts.

Our results also show that to evidence Fitts' law a multi-scale Fitts task need not involve a terminal target *selection* act (e.g., a mouse click). Our participants had to reach a target item of size W at a distance D , which allowed us to calculate an ID . However, our Fitts task had one rather unusual feature: participants did not have to select items, they were simply to visualize them (by scrolling and rescaling) so as to make them readable. The double click they were to produce for each target was spatially unconstrained—obviously, a vocal response would have done just as well in our experiment. Such a task differs notably from the standard Fitts task, yet we obtained evidence that target-acquisition time is still linearly dependent on Fitts' ID (Fig. 9). This broadens the usability of target-acquisition time for experiments on multi-scale navigation: Fitts' law accounts for the multi-scale *visualization* of target items, regardless of whether or not a final pointing act is required. In this sense the scope of Fitts' ID for evaluation studies is probably wider than has been thought so far in HCI.

Turning to our second intention—explaining the benefits of a new technique—it should be clear from the preceding that

we are *not* advocating PD as an alternative to the usual Z&D technique. What we do claim is that for everyday document navigation it would help users to have at their disposal one degree of freedom of camera tilt, in addition to the three degrees of freedom of translation to which they have been currently entitled so far. Indeed, depending on the nature of the navigation task they face, users should be allowed to choose the navigation technique they feel is most convenient to them.

With regard to the value of PD, what was learned from the above two experiments can be summarized in either of two ways. We may say that the more diluted the payoff of reaching the terminal goal, the more useful PD, or equivalently that the more difficult the aimed navigation (in the cumulated sense expressed by the ID_c), the more useful the PD technique. Either way, there is serious reason to consider that novel technique for implementation as a default resource in GUIs.

It seems quite likely that navigation tasks with a moderate level of DGD and, accordingly, a fairly high level of cumulated difficulty form a large proportion of what users actually do in offices and homes. This, we believe, is a strong argument for promoting techniques that work best with moderately difficult tasks, as is the case with PD.

ACKNOWLEDGMENTS

This research was supported by the French *Ministère Délégué à la Recherche et aux Nouvelles Technologies (ACI Masses de données, Micromegas project)* granted to the first author. We thank C. Appert, M. Beaudouin-Lafon, S. Benhamou, and E. Pietriga for stimulating discussions, and N. Roussel for not only sharing ideas, but also developing our very first research prototype.

REFERENCES

- Bourgeois, F., Guiard, Y., and Beaudouin-Lafon, M. Pan-zoom coordination in multi-scale pointing. *Ext. Abstracts CHI 2001*, ACM Press (2001), 157-158.
- Card, S.K., English, W.K., and Burr, B.J. Evaluation of mouse, rate-controlled isometric joystick, step-keys, and text keys for text selection on a CRT. *Ergonomics* 21 (1978), 301-613.
- Fitts, P.M. The information capacity of the human motor system in controlling the amplitude of movement. *J. Exp. Psychology*. 47 (1954), 381-391. Reprinted in *J. Exp. Psychology: General* 121, 3 (1992), 262-269.
- Foley, J.D., van Dam, A., Feiner, S.K., and Hughes, J.F. *Computer Graphics: Principles and Practice*. Addison-Wesley Longman, Boston, MA, USA, 1990 (2nd Ed.).
- Furnas, G. Generalized fisheye views. *Proc. CHI 1986*, ACM Press (1986), 16-23.
- Furnas, G. A fisheye follow-up: Further reflections on focus + context. *Proc. CHI 2006*, ACM Press (2006), 999-1008.
- Gibson, J.J. *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston, MA, USA, 1979.
- Guiard, Y., and Beaudouin-Lafon, M. Target acquisition in multi-scale electronic worlds. *Internat. J. Human-Computer Studies* 61 (2004), 875-905.
- Guiard, Y., Bourgeois, F., Mottet, D., and Beaudouin-Lafon, M. Beyond the 10-bit barrier : Fitts' law in multi-scale electronic worlds. *Proc. IHM-HCI 2001*, Springer Verlag (2001), 573-587.
- Guiard, Y., Chapuis, O., Du, Y., and Beaudouin-Lafon, M. Allowing camera tilts for document navigation in the standard GUI: A discussion and an experiment. *Proc. AVI'06*, ACM Press (2006), 241-244.
- Lee, J.P., and Grinstein, G. An architecture for retaining and analyzing visual explorations of databases. *Proc. VIS 1995*, IEEE Computer Society (1995), 101-108.
- Mackay, W.E. Which interaction technique works when? Floating palettes, marking menus and toolglasses support different task strategies. *Proc. AVI 2002*, ACM Press (2002), 203-208.
- MacKenzie, I.S. Fitts' law as a research and design tool in human-computer interaction. *Human-Computer Interaction* 7 (1992), 91-139.
- Mackinlay, J.D., Robertson, G.G., and Card, S.K. The perspective wall: detail and context smoothly integrated. *Proc. CHI 1991*, ACM Press (1991), 173-176.
- Pietriga, E., Appert, C., and Beaudouin-Lafon, M. Pointing and beyond: an operationalization and preliminary evaluation of multi-scale searching. *Proc. CHI 2007*, ACM Press (this volume).
- Pirolli, P., and Card, S. Information foraging in information access environments. *Proc. CHI 1995*, ACM Press (1995), 51-58.
- Sen, A. *Inequality reexamined*. Oxford University Press, Oxford, UK, 1992.
- Soukoreff, R.W., and MacKenzie, I.S. Towards a standard for pointing device evaluation: Perspectives on 27 years of Fitts' law research in HCI. *International J. Human-Computer Studies* 61 (2004), 751-789.
- Wiltschko, R., and Wiltschko, W. *Magnetic Orientation in Animals*. Springer-Verlag, Eidelberg, Germany, 1995.