Table 2 summarizes the viewpoint controls of Figure 1 associated with the above four tasks to be supported by our particular application. Having identified the viewpoint model's degrees of freedom we want to control and taken a few first steps into their organization, we must now turn to the input device (the multi-touch screen) and examine the handles it provides for that control.

	T_x	T_y	T_x	R_x	R_y	FOV	O_y
Move around	◐	◐	●	○	●	○	○
Look around	○	○	○	●	●	○	○
Circle around *P*	○	○	○	○	○	○	●
Scrutinize *P*	○	○	○	○	○	●	○

Table 2. Relevance of viewpoint controls to high-level navigation tasks, by decreasing order of importance.●: relevant,◐: partially relevant, ○: not relevant.

3.2 Identifying input handles

Although some touch sensing technologies provide rich information about contact regions, including their shape or the applied force, most multi-touch APIs only expose the 2D coordinates of their centroid. One might thus think that using n fingers, users should be able to control $2 \times n$ degrees of freedom. But in reality, it is never the case. Multi-touch systems can't distinguish between fingers, so degrees of freedom cannot be univocally associated to them. The order of appearance of contacts or hit-testing with specific on-screen areas can be used for these associations. But in the end, interaction will always be constrained by limited finger individuation: it is quite difficult to move one finger without some degree of involuntary movement at one or more of the others [11].

Instead of considering contacts individually, one can group them using different methods (e.g. hit-testing, proximity, hand identification) and extract from the collated movement information global parameters to be associated with degrees of freedom to control. A common practice is to consider multi-touch gestures on objects as *Rotate-Scale-Translate* (RST) manipulations and to determine and characterize the principal transformation involved - e.g. $R(\alpha, C)$ for a *turn* gesture, $S(\kappa, C)$ for a *pinch* or a *spread*, and $T(x, y)$ for a *swipe*[2]. Figure 3 shows a simplified view of the state machine we used, based on this approach. The machine differentiates four interaction states (shown in gray): one for single-touch interaction, and three differentiating multi-touch interactions based on the first principal transformation detected.

Having described the desirable viewpoint controls for our application (Table 2) and the different handles provided by a multi-touch screen, i.e. $T(x, y)$ for single-touch interactions and $R(\alpha, C)$, $S(\kappa, C)$ & $T(x, y)$ for multi-touch ones, we will now examine the mappings between them.

[2] The parameters associated to each transformation will be explained in the next section.

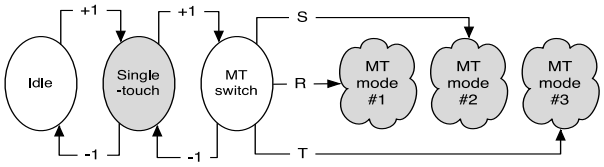

Figure 3. Simplified view of a multi-touch state machine based on an RST classifier. +1 and -1 transitions are triggered by contact addition and removal. The R, S and T transitions are triggered based on the first principal transformation detected. Note that once in mode #1, #2 or #3, multi-touch interaction does not have to be restricted to the transformation initially detected.

3.3 Choosing the right mappings

A way to reduce the articulatory distance and ensure tight coupling between perception and action is to choose the mappings between viewpoint controls and multi-touch gestures so that contact trajectories match the scene transformations caused by viewpoint modifications. To achieve this, we created video clips illustrating the effect of the 7 viewpoint controls of Table 2 on a particular scene. We used OpenCV to compute and visualize the optical flow of each of these videos (Figure 4). We then compared these flows to those corresponding to multi-touch RST manipulations, i.e. $T(x, y)$ when using a single finger and $R(\alpha, C)$, $S(\kappa, C)$, $T(x, y)$ otherwise.

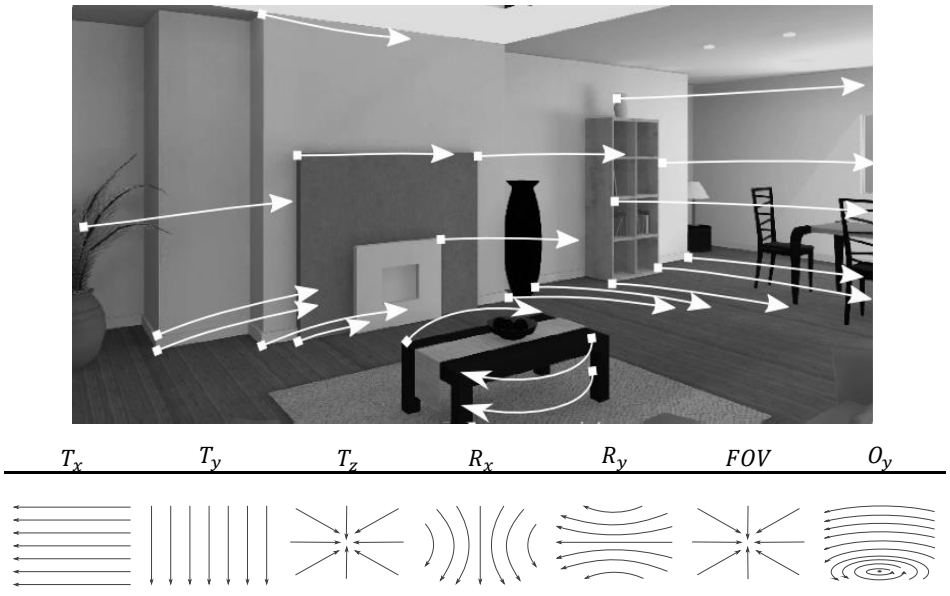

Figure 4. Top image: optical flow computed over time with OpenCV when circling the viewpoint along the vertical axis (R_y). Bottom images: stylized renderings of the flows corresponding to the 7 viewpoint controls of Table 2 (arrows could all point in the opposite direction).

Table 3 summarizes the compatibility between gesture flows and optical flows using a three-level scale. Starting from this table, we applied the following heuristics to choose between alternative mappings:

- As MT-$R(\alpha, C)$ is the most compatible gesture with O_y, we decided to map the two.
- T_z and FOV are compatible with the same gestures, MT-$S(\kappa, C)$ and ST-$T(x, y)$. Since T_z is one of the most important controls, we wanted to keep it as simple as possible and thus preferred a single-touch gesture for it. Moving forward/backward seemed better matched with a vertical movement rather than a horizontal one, so we chose ST-$T(., y)$. For FOV, we chose MT-$S(\kappa, C)$.
- For ST-$T(x, .)$, we were left with T_x, R_y and O_y, the first two being more important than the third one. We decided to map ST-$T(x, .)$ to R_y so that single-touch interaction would support both *Move around* (with T_z and R_y) and *Look around* (with R_y).
- For MT-$T(., y)$, we were left with T_y and R_x. We chose the latter, as looking up/down was considered more important than controlling one's altitude.
- For MT-$T(x, .)$, we were left with T_x and R_y. Informal tests convinced us that the latter was preferable, considering our previous choice of R_x for MT-$T(., y)$.

		T_x	T_y	T_z	R_x	R_y	FOV	O_y
ST	$T(x, .)$	●	○	●	○	◐(1)	●	◐(2)
	$T(., y)$	○	●	●	◐(3)	○	●	○
MT	$R(\alpha, C)$	○	○	○	◐(4)	◐(5)	○	●
MT	$S(\kappa, C)$	○	○	●	○	○	●	○
MT	$T(x, .)$	●	○	◐(6)	○	◐(1)	◐(6)	◐(2)
	$T(., y)$	○	●	◐(6)	◐(3)	○	◐(6)	○

Table 3. Compatibility between gesture flows (rows) and optical flows (columns): ○ incompatible, ● compatible, ◐ compatible under one of the conditions below. A dot in place of x or y indicates that this component is ignored. The gray cells correspond to the chosen mapping.

(1) compatibility is inversely proportional to the vertical distance to the center of the screen
(2) the point of interest (P) must have been previously specified
(3) compatibility is inversely proportional to the horizontal distance to the center of the screen
(4) C must be in the middle of a vertical border of the screen, i.e. left or right
(5) C must be in the middle of a horizontal border of the screen, i.e. top or bottom
(6) contacts must be "close enough", i.e. within a certain radius, so they can be reduced to single-touch interaction

4 IMPLEMENTATION: THE MOVE&LOOK TECHNIQUE

A close look at Table 2 and the chosen mapping in Table 3 shows that ST-$T(x, y)$ corresponds to *Move around* while MT-$R(\alpha, C)$ corresponds to *Circle around*, MT-$S(\kappa, C)$ to *Scrutinize* and MT-$T(x, y)$ to *Look around*. As illustrated by Figure 5, each of our 4 high-level navigation tasks can thus be associated to an interaction state of the machine shown in Figure 3. This section details the implementation of the resulting navigation technique, which we called *Move&Look*.

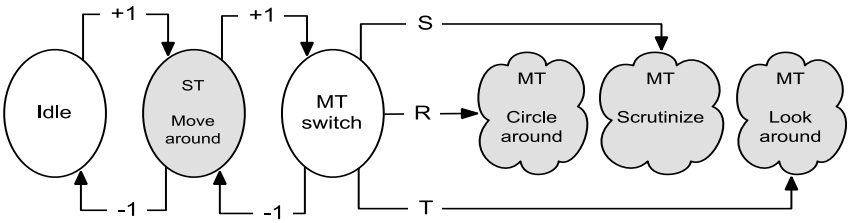

Figure 5. The Move&Look technique, instanciated from Figure 3.

4.1 Single-touch interaction: *move around*

When a single contact is detected, subsequent $T(x, y)$ finger movements are mapped to R_y, T_z camera movements to support the *move around* task. When touch is detected, a ray is casted in the 3D scene from the camera center through the contact point in the camera plane. The intersection with the scene (P) defines the point of interest, and the ray a path towards it. Progression along the path is controlled through finger movements in the following way:

- Lateral movements ($T(x, .)$) do not affect the camera position. Proximal finger movements translate the camera towards P and distal movements translate it backwards along the path ($T(., y)$).
- The distance between the initial contact point and the bottom of the display is mapped to the entire path length: the destination is reached when the finger reaches the bottom of the display.
- Distal finger movements past the initial contact point (i.e. above it) continue moving the camera backwards along the path. For consistency, the scale factor remains the same as when closing in on P.

Users can turn the camera left and right (R_y) through lateral finger movements. The camera orientation is computed so as to always keep the projection of P under the finger. R_y is computed either analytically [7] or numerically [23] to minimize the distance between the previous projection of P and the current finger position (we used the minimizer from ALGLIB to solve the different minimization problems).

4.2 Multi-touch interaction switch: RST classifier

When multiple contacts are detected, their movement is analyzed to determine whether the state machine should switch to *circle around*, *scrutinize* or *look around*. The movement of the n contact points is interpreted as a rigid transformation combining a rotation $R(\alpha, C)$, an homogeneous scaling $S(\kappa, C)$, and a translation $T(x, y)$. The initial position of the contact points is noted r_i and their current position c_i. The R, S and T transformations correspond to the minimization of the cost function F defined by Equation 1. $T(x, y)$ is first computed from the centroids of the initial set of contact points (C_i) and the current one (C_c) according to equations 2, 3 and 4. The rotation angle α is the one that minimizes the cost function of Equation 5 and is computed using Equation 6 where $co = c_i - C_i$ and $ro = r_i - C_c$. The scale factor κ is similarly

the one that minimizes the cost function of Equation 7 and is computed using Equation 8 where $cr = R^{-1}(\alpha).c$:

$$F(x,y,\alpha,\kappa) = \sum_{0\le i\le n} \|T(x,y)R(\alpha)S(\kappa)r_i - c_i\| \quad (1)$$

$$C_i = \frac{1}{n}\sum_{0\le i\le n} r_i \quad (2)$$

$$C_c = \frac{1}{n}\sum_{0\le i\le n} \mathrm{c_i} \quad (3)$$

$$T(x,y) = C_c - C_i \quad (4)$$

$$G(\alpha) = \sum_{0\le i\le n} \|R(\alpha)(r_i - C_i) - (c_i - C_c)\| \quad (5)$$

$$\alpha = \operatorname{atan}2\left(-\sum_{0\le i\le n} co_x ro_y - co_y ro_x\,,\ \sum_{0\le i\le n} co_y ro_y + co_x ro_x\right) \quad (6)$$

$$H(\kappa) = \sum_{0\le i\le n} \|S(\kappa)ro - R^{-1}(\alpha)co\| \quad (7)$$

$$\kappa = \frac{cr_x^2 + cr_y^2}{cr_x ro_x + cr_y ro_y} \quad (8)$$

The only rotations considered in these equations are those centered on the centroid of the contact points. Although we perceive it as an elementary rotation, moving one's index finger around one's thumb while keeping this one steady will thus be interpreted as a combination of a centroid-centered rotation and a translation. To tackle this problem, we weight all contact positions by the inverse of their traveled distance when computing C, the center of both the rotation $R(\alpha, C)$ and the homogeneous scaling $S(\kappa, C).T(x,y)$ is also adjusted by removing the displacement possibly introduced by $R(\alpha, C)$.

The x, y, C, α and κ parameters resulting from the above computations are used to determine the prominent gesture among *Rotate*, *Swipe* and *Pinch*. Our classifier considers one model $M(t)$ for each gesture type and returns the one that better fits the observations (highest R^2 value). The models map the initial configuration r_i to an estimated state $\widetilde{c_i}$ (Equation 9). The estimated error (the residual sum of squares) and the coefficient of determination R^2 are then computed using Equations 10, 11 and 12.

$$\widetilde{c_i} = M(t).r_i \quad (9) \qquad R^2 = 1 - \frac{SS_{res}}{SS_{tot}} \quad (10)$$

$$SS_{res} = \sum_{0\le i\le n} \|c_i - \widetilde{c_i}\|^2 \quad (11) \qquad SS_{tot} = \sum_{0\le i\le n} \|r_i - c_i\|^2 \quad (12)$$

Our classifying method is similar to the *GestureMatching* method used by Nacenta et al. [21], but instead of classifying the combined contribution of *Rotate*, *Swipe* and *Pinch* gestures, ours allows to classify the contribution of individual gesture types. Our method requires enough information to properly work. The classifier is thus enabled only when the summed distance covered by the contact points is beyond a given threshold. Based on preliminary tests, we found that a value of 10 pixels on a 90PPI display (around 2.8 mm) provides a good trade-off between latency and success rate, which is in agreement with other thresholds reported in the literature for similar applications [25].

4.3 Multi-touch gestures: *circle around*, *look around* and *scrutinize*

When the classifier detects a prominent *Rotate* gesture, the technique enters the *circle around* state of Figure 5 until all contacts are lost. $R(\alpha, C)$ provides the pivot point to rotate the scene and the angle of rotation (the center of the 3D rotation is computed from the projection of C in the 3D scene).

When a *Swipe* gesture is detected, the technique enters the *look-around* state until all contacts are lost. $T(x, y)$ is then used to rotate the camera (R_x, R_y) in a way similar to *move around*, but with two degrees of freedom instead of one.

When a *Pinch* gesture is detected, the technique enters the *scrutinize* state until all contacts are lost. The scale factor of $S(\kappa, C)$ is used to adjust the FOV of the camera. To ensure smooth camera movements, its look-at point remains fixed while contacts are moving. The FOV is restored to its initial value when all contacts are lost. This state further supports remote inspection by using $T(x, y)$ to rotate the camera (R_x, R_y), as in the *look-around* state.

5 EXPERIMENT

Our main motivation in this experiment was to assess our design choices by comparing Move&Look to other techniques from the literature (Screen-space [23], DabR [5] and Drag'n Go [20]) or available in commercial products (Virtual joysticks and the RealMyst technique described above), most of which have never been evaluated nor compared.

5.1 Task

Informal user testing with Move&Look suggested the technique was particularly efficient for interior designs mainly consisting of flat orthogonal surfaces. Encouraged by this, we wanted to assess the effectiveness of the technique in a more demanding environment. The task we chose consisted in collecting spheres placed inside boxes in an outdoor environment, and dropping them in a fountain at the center of the scene (Figure 6). To provide a fair comparison between multi-scale navigation techniques (Drag'n Go and Move&Look) and the others and focus on the evaluation of camera movements, the boxes were not positioned far away from each other but close to the central drop zone.

Participants had to find the boxes in the scene. For each box, they had to position the camera in front of its only open face to pick up the sphere it contained. A sphere

would turn from red to green when the camera was close enough to indicate one could touch it to pick it up. Participants could carry only one sphere at a time, and it was automatically dropped once in the drop zone. Collision detection prevented participants from moving through objects, and a trial was considered as fully completed after all the spheres had been dropped. Participants were instructed to perform this as quickly as possible. They could ask the experimenter to reset the camera to its initial position or withdraw a trial if they felt unable to complete it. They were not encouraged to do so, however. The experimenter rather encouraged them to finish a trial if he felt they could succeed.

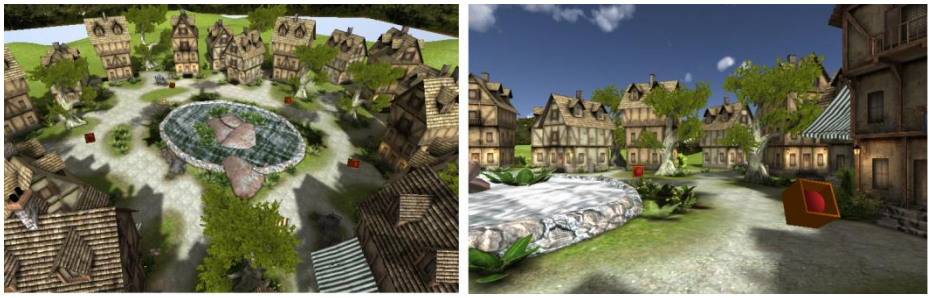

Figure 6. Left: overview of the 3D environment used in the experiment. Right: detailed view showing a box containing a sphere to pick up and drop in the fountain.

5.2 Participants

Twelve unpaid male participants with a mean age of 35 (SD=14) served in the experiment. Five of them used a computer on a daily basis, played 3D video games and were familiar with touch-screens through mobile phones or tablets. Seven of them had a low experience with video games and were novice with touch interfaces.

5.3 Apparatus, design and procedure

Participants were seated in front of a 22" 3M multi-touch screen orientated at an angle of about 70° from a horizontal desk. The experiment was implemented using Unity 3.5[3]. A repeated measures within-subjects design was used. The independent variable was the interaction technique (TECH) with six levels: DabR, Screen-space, Drag'n Go, RealMyst (a custom implementation of the RealMyst technique), Virtual joysticks (a standard combination of two Unity virtual joysticks displayed at fixed positions) and Move&Look. A trial consisted in collecting 4 spheres and each technique was evaluated with 3 successive trials (TRIAL). In summary the experimental design was: 12 participants × 6 TECH × 3 TRIAL = 212 total trials.

The presentation order for TECH was counter-balanced across participants using a balanced Latin Square design. To favor expert usage and a fair comparison between techniques, each was first introduced by the experimenter with a demo and then a training session. Participants could also use a cheat sheet throughout the experiment. After each technique, participants filled a questionnaire inspired by the Nasa TLX test

[3] http://unity3d.com/unity/whats-new/unity-3.5

and at the end of the experiment, they were asked to rank the techniques and give additional feedback.

5.4 Results

The dependent variables were the completion time, the number of give-ups and camera resets, and the qualitative results.

Numbers of give-ups and camera resets — 25% of trials were aborted for Screen-space, 17% for RealMyst, 8% for Virtual joysticks and 0% for Drag'n Go, DabR and Move&Look. The camera was reseted in 66% of all trials for Screen-space, 14% for Virtual joysticks, 11% for RealMyst, 3% for DabR and Move&Look, and 0% for Drag'n Go.

Task completion time — Task completion time is defined as the time needed to successfully collect the four spheres and drop them in the fountain. Trials where participants gave up were removed for the analysis. Trials at least three standard deviations away from the mean for each TECH condition were considered as outliers and also removed. A repeated measures ANOVA showed a significant effect of TECH ($F_{5,55} = 9.9, p < 0.001$). Subsequent pairwise comparison showed significant differences ($p < 0.005$) between Drag'n Go and Screen-space, Drag'n Go and RealMyst, Move&Look and Screen-space, and Move&Look and RealMyst. No significant difference was found between Drag'n Go and Move&Look. Completion times were $97s$ for DabR, $184s$ for Screen-space, $60s$ for Drag'n Go, $117s$ for RealMyst, $111s$ for Virtual joysticks and $72s$ for Move&Look.

User ranking and questionnaire — The participants ranked the techniques in decreasing order of preference. Overall, Move&Look came first (10 participants ranked it first and 2 ranked it second) followed by Drag'n Go, DabR, Virtual joysticks, RealMyst and Screen-space. The participants who ranked Move&Look first explained it nicely complements Drag'n Go as it allows to control more degrees of freedom while keeping the navigation intuitive: it does not require focusing on the gestures to execute nor does it require planning a trajectory in the scene to reach a target. Screen-space was ranked last considering its lack of intuitiveness: in spite of the frequent use of the cheat sheet the participants did not understand how to effectively use the technique to navigate the way they wanted (we believe the important semantic distance explains this gap between users' intentions and the system's behaviours). These subjective results are in agreement with the quantitative results found for completion time and the numbers of give-ups up and camera resets.

After each technique, the participants answered questions related to the following six criteria on a 5 point Likert scale: mental demand, physical demand, performance, effort, frustration and satisfaction. The questions asked were similar to the ones available in the Nasa TLX test. We ran a Friedman analysis with Bonferroni-corrected Wilcoxon post-hoc analyses. This analysis shows significant differences between the techniques for all criteria, especially for the techniques at the bottom of the participants' ranking. Table 5 summarizes the significant differences that were found.

	Mental demand	Physical demand	Perfor-mance	Effort	Frustra-tion	Satisfac-tion
Move&Look-Screen-space	●	●	●	●	●	●
Move&Look-RealMyst	●	●	●	○	●	●
Drag'n Go-Screen-space	●	●	●	●	○	●
Drag'n Go-RealMyst	●	●	●	○	○	○
Drag'n Go-DabR	○	●	○	○	○	○
RealMyst-Virtual joysticks	○	●	○	○	○	○
RealMyst-DabR	○	○	○	○	●	○
Screen-space-DabR	●	○	○	○	●	●
Screen-space-Virtual joystick	●	●	○	○	○	●

Table 5. Details of the post-hoc analysis for cases where one or more significant differences were found (●: significant difference, ○: non significant difference).

User feedback and observations — During the experiment we encouraged the participants to "think aloud" and freely comment on the interaction techniques. Comments were overall in agreement with the user ranking.

Screen-space received the most negative critics. All participants repeatedly reported their frustration with this technique. They felt out of control and found the mappings between fingers and camera movements inconsistent. The finger movements corresponding to different screen-space controls can indeed be quite similar, as illustrated by the optical flows of Figure 4 (e.g. T_x and R_y). The output of the screen-space solver is also strongly influenced by the picked point in the 3D scene, and thus by the geometrical shape of the underlying objects. Lastly, the movements to execute in order to move forward (T_z) and to zoom (FOV) depend on whether the initial contact point is above or below the invisible horizon (in the former case, one has to move up, in the latter, one has to move down). All these reasons probably contribute to the fact that users were not able to anticipate camera motions. The comparison of Screen-space to other interaction techniques in 3D manipulation tasks corroborates these observations [18].

Participants found DabR, Virtual joysticks and RealMyst either too slow or too fast. We hypothesize this was caused by the use of transfer functions not specifically tuned for the particular 3D environment we used: long distances took too much time to travel while participants traveled too fast on short distances. Participants found the Virtual joysticks to be less fatiguing. We hypothesize this was due to the use of rate control, which reduces physical movements. Participants reported an important fatigue when using DabR and complained they had to pay attention to the number of fingers they used. They complained about the delay introduced by the time-based mode switch used by RealMyst and the fact that the traveling direction is not towards the selected point but along the T_z axis of the camera. Drag'n Go was particularly appreciated for its ability to quickly reach distant targets, but moving to a box while orienting the viewpoint in order to pick the sphere was found more difficult and required some planning. This was not reported as a problem with Move&Look.

6 Conclusion and future work

In this paper, we proposed an original methodology based on user-centered practices and optical flow analysis to address the problem of designing intuitive multi-touch navigation techniques for 3D environments. User-centered practices allow to define the navigation commands from the user's perspective while the optical flow analysis provides guidelines for defining intuitive multi-touch gestures to perform these commands. We instantiated this methodology for tasks articulated around the review of interior designs, which led to the design of a new interaction technique, Move&Look. The comparison of this technique to state of the art ones in a controlled experiment showed its overall superiority and revealed usability problems with the others. These results provide a first validation of the proposed design methodology. The methodology should be applied in other navigation contexts in order to further assess its effectiveness. The robustness of the proposed RST classifier should be formally evaluated, and it can certainly be improved. Even if participants did not complain about it, we observed them flattening their rotation gestures for the *circle around* command, probably because they unconsciously followed the corresponding optical flow. Our classifier could be modified to better take into account this oval shape, for example.

References

1. Ajaj, R., Vernier, F., Jacquemin, C.: Follow my finger navigation. In: Proc. of INTERACT '09. pp. 228-231. Springer - Verlag (2009).
2. Bowman, D.A., Kruijff, E., LaViola, J.J., Poupyrev, I.: 3D user interfaces: theory and practice. Addison-Wesley/Pearson Education (2003).
3. Britton, E.G., Lipscomb, J.S., Pique, M.E.: Making nested rotations convenient for the user. In: Proc. of SIGGRAPH '78. pp. 222 - 227. ACM (1978).
4. Christie, M., Olivier, P.: Camera control in computer graphics: models, techniques and applications. In: ACM SIGGRAPH ASIA 2009 Courses. pp. 3:1 - 3:197. ACM (2009).
5. Edelmann, J., Schilling, A., Fleck, S.: The DabR - a multitouch system for intuitive 3D scene navigation. In: Proc. of 3DTV-CON. pp.1 - 4. (2009).
6. Fu, C.W., Goh, W.B., Ng, J.A.: Multi-touch techniques for exploring large-scale 3D astrophysical simulations. In: Proc. of CHI '10. pp. 2213 - 2222. ACM (2010).
7. Gleicher, M., Witkin, A.: Through-the-lens camera control. Proc. of SIGGRAPH '92 26, 331 - 340. (July 1992).
8. Hachet, M., Decle, F., Knodel, S., Guitton, P.: Navidget for easy 3D camera positioning from 2D inputs. In: Proc. of 3DUI '08. pp. 83 - 89. IEEE (2008).
9. Hutchins, E.L., Hollan, J.D., Norman, D.A.: Direct manipulation interfaces. Human-Computer Interaction 1(4), 311 - 338. (Dec 1985).
10. Ingram, A., Wang, X., Ribarsky, W.: Towards the establishment of a framework for intuitive multi-touch interaction design. In: Proc. of AVI '12. pp. 66 - 73. ACM (2012).

11. Ingram, J.N., Körding, K.P., Howard, I.S., Wolpert, D.M.: The statistics of natural hand movements. Experimental brain research 188(2), 223 - 236. (Jun 2008).

12. Jacob, R.J.K., Sibert, L.E., McFarlane, D.C., Mullen, Jr., M.P.: Integrality and separability of input devices. ACM ToCHI 1(1), 3 - 26. (Mar 1994).

13. Jacob, R.J., Girouard, A., Hirshfield, L.M., Horn, M.S., Shaer, O., Solovey, E.T., Zigelbaum, J.: Reality-based interaction: a framework for post-wimp interfaces. In: Proc. of CHI '08. pp. 201 - 210. ACM (2008).

14. Kim, J.S., Gračanin, D., Matković, K., Quek, F.: iPhone/iPod Touch as input devices for navigation in immersive virtual environments. In Proc. of VR '09. 261 - 262. IEEE.

15. Kim, J.S., Gračanin, D., Matković, K., Quek, F.: Finger Walking in Place (FWIP): A traveling technique in virtual environments. In Proc of SG '08 58 - 69. Springer-Verlag.

16. Kruger, R., Carpendale, S., Scott, S.D., Tang, A.: Fluid integration of rotation and translation. In: Proc. of CHI '05. pp. 601 - 610. ACM (2005).

17. Mackinlay, J.D., Card, S.K., Robertson, G.G.: Rapid controlled movement through a virtual 3D workspace. In: Proc. of SIGGRAPH '90. pp. 171 - 176. ACM (1990).

18. Martinet, A., Casiez, G., Grisoni, L.: Integrality and separability of multitouch interaction techniques in 3D manipulation tasks. IEEE TVCG 18(3), 369 - 380. (Mar 2012).

19. McCrae, J., Mordatch, I., Glueck, M., Khan, A.: Multiscale 3D navigation. In: Proc. of I3D '09. pp. 7 - 14. ACM (2009).

20. Moerman, C., Marchal, D., Grisoni, L.: Drag'n Go: Simple and fast navigation in virtual environment. In: Proc. of 3DUI '12. pp. 15 - 18. IEEE (2012).

21. Nacenta, M.A., Baudisch, P., Benko, H., Wilson, A.: Separability of spatial manipulations in multi-touch interfaces. In: Proc. of GI '09. pp. 175 - 182. Canadian Information Processing Society (2009).

22. Rasmussen, J.: Skills, rules, and knowledge; signals, signs, and symbols, and other distinctions in human performance models. IEEE Transactions on Systems, Man and Cybernetics 13(3), 257 - 266. (May 1983).

23. Reisman, J.L., Davidson, P.L., Han, J.Y.: A screen-space formulation for 2D and 3D direct manipulation. In: Proc. of UIST '09. pp. 69 - 78. ACM (2009).

24. de la Rivière, J.B., Kervégant, C., Orvain, E., Dittlo, N.: CubTile: a multi-touch cubic interface. In: Proc. of VRST '08. pp. 69 - 72. ACM (2008).

25. Zeleznik, R., Forsberg, A.: UniCam - 2D gestural camera controls for 3D environments. In: Proc. of I3D '99. pp. 169 - 173. ACM (1999).