# The Answer Is Blowing in the Wind: Directed Air Flow for Socially-acceptable Human-Robot Interaction

Vincent Zhang[1][a], Natalie Friedman[1,2][b], David Goedicke[1,2][c], Dmitriy Rivkin[1][d], Michael Jenkin[1,3][e], Xue Liu[1,4][f] and Gregory Dudek[1,4][g]

[1]*Samsung AI Center, Montreal, Canada*

[2]*Information Science, Cornell Tech, New York, U.S.A.*

[3]*EECS, York University, Toronto, Canada*

[4]*Computer Science, McGill University, Montreal, Canada*

Keywords: Human-Robot Interaction, Air-driven Haptic Cues.

Abstract: A key problem for a robot moving within a social environment is the need to capture the attention of other people using the space. In most use cases, this capture of attention needs to be accomplished in a socially acceptable manner without loud noises or physical contact. Although there are many communication mechanisms that might be used to signal the need for a person's attention, one particular modality that has received little interest from the robotics community is the use of controlled air as a haptic signal. Recent work has demonstrated that controlled air can provide a useful signal in the social robot domain, but what is the best mechanism to provide this signal? Here, we evaluate a number of different mechanisms that can provide this attention-seeking communication. We demonstrate that many different simple haptic air delivery systems can be effective and show that air on and air off haptic events have very similar time courses using these delivery systems.



## 1 INTRODUCTION

The real world presents many challenges to a mobile robot that are not encountered in a controlled laboratory setting. In particular, in the real world the motion plans of mobile robots are often obstructed by humans and getting humans to move out of the way can be challenging. Autonomous vehicles need mechanisms to attract the attention of humans, perceived as, "obstacles" in a socially acceptable manner. Classic methods that are often found in industrial settings, like beeping loudly or waiting patiently, are either disruptive or ineffective (Wogollter, 2006) and sometimes both when used in a social setting.

A classic example of this problem is the motion of a robot through a crowd in a social setting as shown in Figure 1. For the robot to be able to make its way to the goal, it must acquire the attention of the people who are obstructing its path and then communicate to


Figure 1: A telepresence robot trying to make its way through a group of individuals who are blocking its path, but is blocked.

[a] https://orcid.org/0000-0003-1134-045X
[b] https://orcid.org/0000-0003-4751-7739
[c] https://orcid.org/0000-0002-4837-893X
[d] https://orcid.org/0000-0003-4136-4831
[e] https://orcid.org/0000-0002-2969-0012
[f] https://orcid.org/0000-0001-5252-3442
[g] https://orcid.org/0000-0001-5040-4925



Zhang, V., Friedman, N., Goedicke, D., Rivkin, D., Jenkin, M., Liu, X. and Dudek, G.
The Answer Is Blowing in the Wind: Directed Air Flow for Socially-acceptable Human-Robot Interaction.
DOI: 10.5220/0010136001060113
In *Proceedings of the International Conference on Robotics, Computer Vision and Intelligent Systems (ROBOVIS 2020)*, pages 106-113
ISBN: 978-989-758-479-4

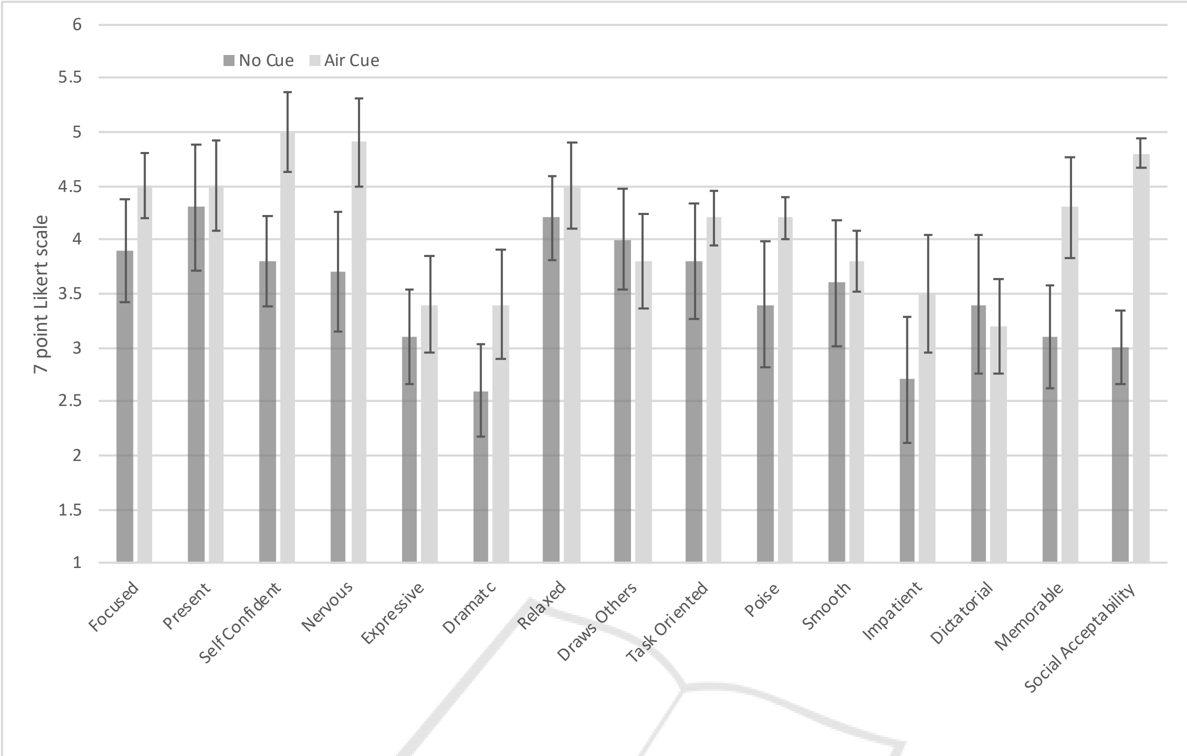


Figure 2: Results from a Wizard of Oz (WoZ) user study on the effect of using haptic air for robot to human interaction (Friedman et al., 2020b). A questionnaire measured the appropriateness of the robot on a 7 point Likert scale. A score of 1 = less focused, less present, less self-confident, etc. A score of 7 = more focused, more present, more self-confident, etc. Error bars show standard errors. Haptic air cues made the robot seems more socially acceptable compared to a robot that did not use this cue over a range of different measures. Results shown here are for the 23 adults (4 female) participants who completed the questionnaire associated with the study. The mean age of the study participants was 25 (SD=5.8). Full details of the study can be found in (Friedman et al., 2020b).

them its intention to move through the space. Capturing the attention of users could be accomplished in a number of ways, including through the use of acoustic, visual and haptic cues in various combinations. But not all strategies are socially acceptable. For example, loud noises are generally unacceptable in a social setting and visual cues are ineffective except when the robot is in direct line of sight. Cues based on direct physical contact are likely to be inappropriate in many social contexts. Touching or tapping is generally a questionable behavior with strangers, especially in some cultures where physical contact is frowned upon. Even the nature of acceptable touching varies from culture to culture. For example, in some cultures, it is acceptable to touch a stranger with an open palm, but potentially rude to do it with a finger tip. In the context of a robot, it is unknown if tapping a person on the shoulder would be a socially appropriate strategy for an autonomous robot.

One potential alternative to an actual physical touch is to use an alternative medium such as directed air as a conduit for this haptic touch. In Friedman et al. (2020a, b) a haptic air cue was integrated with an audio cue and visual instructions in a robot interaction experiment. This Wizard of Oz (WoZ) “in the wild” experiment was centred around a robot attempting to encourage participants in the space to move out of the way as the robot followed a pre-computed path. This study explored the use of air-based haptic cues combined with audio and visual cues to both attract attention and then to communicate intent to participants blocking the robot’s path. Providing the robot with a mechanism to attract attention (haptic air and audio cues) coupled with explicit instruction to the user to move out of the way was found to be an effective strategy to enable the robot to move through the crowd. In this study participants also completed a questionnaire that asked on a 7-point Likert scale (Burgoon et al., 1998) questions related to the observer’s perception of the robot’s behaviour. Figure 2 plots participant responses from the study where a score of one indicates that the participant found the robot to be less focused, less present, less self-confident, etc., while a score of seven indicates that the participant found the robot to

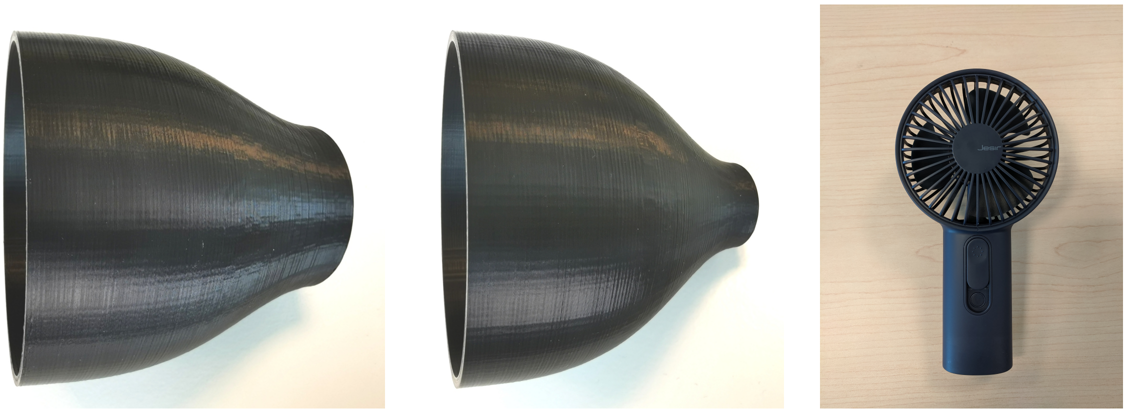

(a) Funnel 1 (3D Printed) (b) Funnel 2 (3D Printed) (c) Common fan

Figure 3: Fan designs. (a) and (b) show the 3D printed cowls of Funnel 1 and Funnel 2. These are mounted to the basic fan shown in (e).

be more focused, more present, more self-confident, and so on. Each of the conditions showed a consistent trend over the majority of the conditions with the addition of air there is a trend for the robot to appear more focused, more present, etc.

The work reported in Friedman (2020a, b) used a very simple fan (shown in Figure 3(c)) to generate the haptic air cue. This choice of fan design was one of convenience. But is this the best mechanism for providing the most appropriate mechanism to generate a haptic air cue? Here we explore a range of technologies that might be exploited in generating this cue.

## 2 RELATED WORK

The use of controlled air has been explored as an interaction medium in the virtual reality/haptics space in a number of different domains and applications. Specific examples include the use of ultrasound transducers (Hoshi et al., 2009; Hasegawa et al., 2017) and arrays of air jets (Suzuki and Kobayashi, 2005) to generate local haptic events. With the goal of simulating large-scale wind sensations, there have also been a number of efforts to build simulated wind environments including Mowafi et al. (2015), Tolley et al. (2019) and Verlinden et al. (2013). An air puff has also been considered as a stimulus in other contexts, for example to elicit responses from children suffering from autism (Dakopolos and Jahromi, 2019). While fans and compressed air are the most commonplace technologies for creating controlled air flows, they are not the only options. Alternative approaches to systems based on mechanical fans and compressed air exist as well. For example Sodhi et al. (2013) and Gupta et al. (2013) describe approaches that use audio subwoofer speakers and an air chamber to generate ring vortex air patterns that are stable and have been found to be capable of being directed out to 1.25m from the emitter.

Although there exist a range of different technologies that might be exploited to generate a controlled haptic air event, perhaps the most straightforward is to use an electric fan coupled with some sort of focusing and aiming device to generate the haptic cue. Even within this controlled design space there are a range of different questions that must be addressed. In particular, are there advantages in terms of shaping the structure of the air column generated by a fan in terms of presenting the cue to the observer? This is the question we consider here.

## 3 FUNNEL DESIGNS AND EVALUATION

An effective controlled air source should provide a constant, non-turbulent flow at a distance of at least a meter to the human participant. Ensuring that the flow is non-turbulent allows the flow to have a non-chaotic structure when it strikes the participant. While a one meter distance allows the robot to provide a haptic cue without entering the personal space of the user but at the same time being sufficiently close to the user that origin of the haptic air event is clear. Furthermore, the width of the haptic air cue should be sufficiently narrow that it can be targeted at the person or persons blocking the robot's path. Narrower cues allows selectivity in terms of the target participant. Clearly, the desired width of the air stream depends on the number of people that the robot needs to contact simultaneously. However, we believe that a

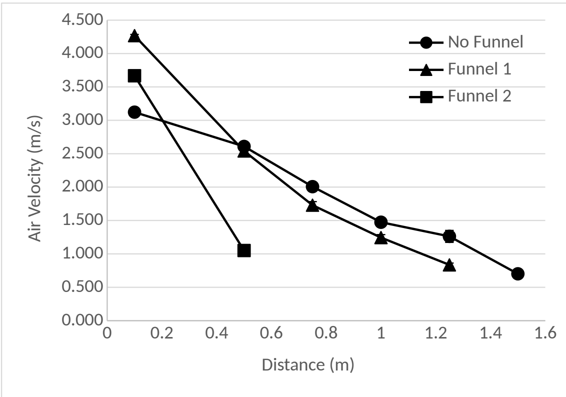


Figure 4: Generated air velocity measured at different distances for the haptic air device designs shown in Figure 3. Air velocities were measured using an anemometer (AP-846A AOPUTTRIVER) at distances of 0.1m to 1.5m from the driving fan.

narrower stream is preferable over a wider one, since it provides the robot the option of alerting a single individual. Groups of people may still be contacted by moving the air stream from side to side. Given these observations, here we evaluate the efficacy of mechanical structures (funnels) in improving the performance of a fixed-speed fan in terms of:

1. The **width** of the air stream in which a lower width means higher performance.
2. The maximum human detection **distance** in which a larger distance means higher performance.
3. Human **reaction time** to “air on” and “air off” events in which shorter is better.

For a given air source (the fan) a funnel mounted on the air source may increase performance in terms of condition 1 (width) by blocking air flow at high off-normal angles. Furthermore, it may also improve performance under conditions 2 (detection distance) and 3 (reaction time) by increasing wind velocity. In order to explore this hypothesis, a number of different funnel structures were constructed and evaluated using the three criteria given above. These funnel designs were circularly symmetric and the radius described by two cubic polynomials, with continuous curvature at the matching point and zero first order derivatives at the inlet and outlet. Several researchers have suggested that this design maximizes air velocity at the outlet by minimizing the adverse pressure gradient along the walls of the funnel (Morel, 1975; Cattafesta et al., 2010). Combining the aforementioned constraint with a length ratio, contraction ratio and match point location fully defines the shape of the funnel.

Figure 5: A view from above of the experimental setup for measuring the width of the air stream. Orange paper streamers are suspended from a bar and viewed from above. Detected deflection of the streamers indicates the presence of an air signal. The effective width of the air signal and its fall off with width can be estimated from the change in deflection as a function of horizontal distance.

In our case, the diameter of the inlet is determined by the size of the fan (105 mm). As shown is Fig. 3, both Funnel 1 and Funnel 2 have the same length ratio as 1:1 and match point $x_m/L_c$ as 1/2. They only differ in terms of contraction ratio. The contraction ratio of Funnel 1 is 4:1, while Funnel 2 has a higher contraction ratio of 16:1. Both funnels were 3D printed using ABS with a layer height of 100 $\mu$m.

## 4 MEASUREMENT OF AIR FLOW

We performed measurements of maximum air speed (Fig. 4) and air stream width (Tab. 1) at several distances from the fan for No Funnel, Funnel 1 and Funnel 2 air sources. An anemometer (AP-846A AOPUTTRIVER) was used to measure air speed. Fig. 4 plots air speed by distance for the various designs. As shown in the figure, an air speed associated with the fan was measurable out to 1.5m. Furthermore, the measured air velocity of Funnel 1 was higher than No Funnel at the exit, but the air flow did not reach as far as the latter. There was a marked decrease in the performance with Funnel 2, which can be explained by reduced air throughput caused by the smaller outlet aperture.

Stream width was measured using a device composed of a contiguous array of thin paper strips. The device was positioned at 0.5m or 1.0m from the fan with the strips aligned orthogonal to the direction of air flow (Fig. 5). The width of the area where strips were lifted by air flow was then measured. Tbl. 1 summarizes the measured air stream widths at 0.5m

Table 1: Measured width of perceptible air flow channel at distance of 0.5m and 1.0m away from the fan.

| Distance | No Funnel | Funnel 1 | Funnel 2 |
|---|---|---|---|
| 0.5m | 49cm | 15cm | 20cm |
| 1.0m | 33cm | 20cm | 0cm |

and 1.0m from the fan. The funnels both produced a significantly narrower air stream at 0.5m than No Funnel. The stream width of Funnel 1 was also far more consistent than No Funnel as the distance was varied. The flow produced by Funnel 2 was not detectable by our device at a distance of 1.0m.

## 5 PERCEPTION OF AIR FLOW

A critical element for any haptic air delivery design is its ability to deliver a perceivable airborne cue at a socially acceptable distance. In other words, "did you feel the robot tap your shoulder" using air flow? While the measurements performed in previous sections can help estimate the range of detectability of a haptic air event, they do not provide much insight about the ability of individuals to note the presence/absence of the haptic air stimulus and the corresponding reaction time. This question was addressed through a user study.

An experiment was performed to identify the maximum distance to which a subject could reliably detect air flow from the various designs. Fig. 6 shows the basic setup. An observer sat in a chair at a measured distance from the haptic air delivery device wearing headphones that played music to obscure any acoustic information from the fan. The fan was left running at all times and the air flow was blocked/released using a large occlusion device. Observers were asked whether they could feel the flow, i.e. if the air stream was flowing or not at a range of distances. For each test case, we used two measurements of 4 (Funnel 2) or 5 (Funnel 1, no funnel) test subjects. Maximum detectability was set to be the last measured distance at which the participant was correct in identifying the state of the haptic air device 80% of the time when they were queried. As shown in Fig. 7, while air flows from No Funnel and Funnel 1 devices could be detected reliably up to 2m, the air from Funnel 2 was only detectable to 1m.

Once the maximum range at which observers could reliably detect the presence/absence of the haptic air stimulus was established, we explored the detectability of "air on" and "air off" events and the delay between the generation of an event (i.e. the insertion or removal of the occlusion) and its perception by the observer. Quantification of this delay is of particular importance in situations where the haptic air stimulus must be synchronized with some other robot action. The same subjects participated in this study as in the sensitivity study reported above. As before, subjects sat in a chair at different distances from the haptic air device with the device positioned to blow air on the back of their head/neck. Each subject completed each condition twice and the mean of their responses was used as their individual response. Subjects wore their normal work clothes and made no particular effort to expose the back of their neck. In other words, everyone's neck was exposed. They were, however, attentive to the stimulus and thus these results should be regarded as the best-case results for detectability; hence the need for the naturalistic user study in the next section.

Figure 6: Quantifying air as a touch cue experimental setup. Observers sat in a chair at a measured distance from the haptic air device and estimated the time it took for them to start and stop feeling the air signal.

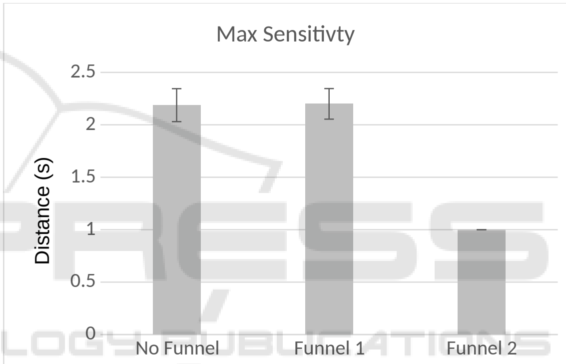


Figure 7: Mean maximum sensitivity distance for the various haptic air delivery device models. Error bars show standard errors.

Figure 8 shows time to detection of the average of *on air* and *off air* events for the No Funnel, Funnel 1 and Funnel 2 designs for different distances. All

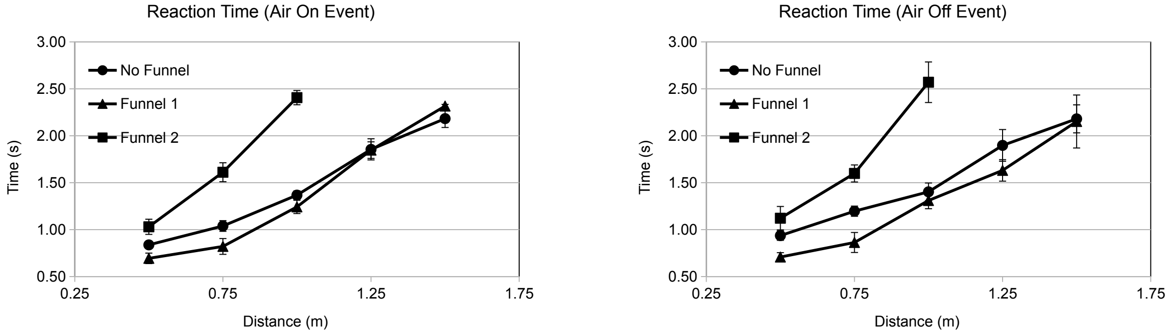


(a) (b)

Figure 8: Time to detection of air on (a) and off (b) events for No Funnel, Funnel 1 and Funnel 2. Error bars show standard errors.

of the curves are quite well fit by straight lines. This suggests that it should be relatively straightforward to properly schedule haptic air on and haptic air off events synchronized with other events on the robot. As shown in Fig. 8, the design of Funnel 1 was able to attract people's attention more effectively than No Funnel at distances within 1.0m, which is in the range of personal distance zone described by Hall (1966), in which someone might tap another on the shoulder. The *air off* event showed a similar time course as the air *on event*, as shown in Fig. 8. Perception of air flow from Funnel 2 was significantly delayed compared with the other two.

One complication with the use of haptic air as a cue is the potential for protective clothing and/or long hair to interfere with the delivery of the haptic cue. In order to consider this we conducted a pilot study with a single individual with long hair using the same procedure as before, but with the hair either being held up or blocking the back of the participant's neck. Fig. 9 shows the sensitivity for this one individual. As expected, having long hair that impedes air flow reduces the efficiency of the cue, but does not eliminate it. For most of the configurations tested, the air cue remained detectable out to 0.75m. Note that this is not the case of Funnel 2 where the subject in the hair down condition was unable to perceive the haptic cue. This suggests that a robot vehicle with an appropriate haptic air generator could potentially infer the appropriate range or air flow for a specific individual based on their neck coverage.

# 6 DISCUSSION

A critical challenge for any robot that is to be deployed in a social setting is maneuvering in coordination with other users in the space. The problem of

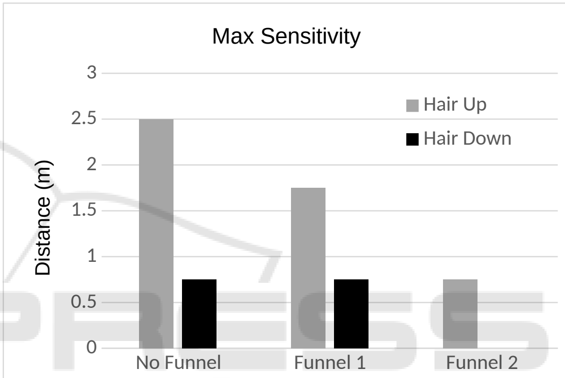


Figure 9: Haptic air sensitivity for one individual with her long hair up and down. Even though there is a considerable decrease in sensitivity, for most of the haptic air devices the cue was perceived reliably out to 0.75m. Note that for the hair down condition for Funnel 2 the subject was unable to perceive the air cue.

moving one agent through a crowd of dynamic obstacles is a classic one, and indeed research in this space can be traced back to Kant and Zucker (1986) . As the space becomes busier, however, it may become impossible to find a path through the space and it then becomes necessary to develop strategies to interact with people in the space to create a path through the space. This requires the robot gaining the attention of people and communicating that the robot needs to get through.

In some domains, the process of gaining the attention of other people in the vicinity can be accomplished by intrusive methods such as loud noises and bright lights. Indeed this is exactly the approach used by emergency vehicles where efficient passage is vastly more important than politeness. In more social settings, such approaches are unlikely to be considered acceptable, hence an autonomous system needs to be able to get the attention of other people in the space in a socially appropriate manner. We have explored the use of fan-generated air flows to acquire this attention and have shown that these flows are capable of delivering a haptic air cue within a range of 2m. We found that the delay between the generation of a haptic air event and a human's perception of it is a predictable function of distance. Therefore, haptic air cuing could be easily integrated into a multimodal attention seeking approach. We also demonstrated that the addition of an appropriately shaped funnel can create a narrower air stream with a more consistent width than a fan alone. A narrow air stream allows the robot to be selective about which humans to contact, allowing it to be less disruptive overall. Funnels can also produce a higher air speed near the fan, resulting in slightly shorter cuing delays at short distances ($<$1m).

Although we explored the use of fans to generate air flow this is not the only potential technology that can be applied to the problem. One promising alternative technology relies on the use of arrays of subwoofer speakers to drive air in a controlled manner. Air cannons driven by such arrays can provide controlled haptic air events. See Sodhi et al. (2013) and Mowafi et al. (2015) for examples of this approach. It is also possible to use stored compressed air to generate easily controlled haptic air events. See Tsalamlal et al. (2014) for an example. Another promising approach is the use of a haptic ultrasound (Long et al., 2014).

In a related study Friedman (2020a, b) we integrated the haptic air cue with an audio cue and visual instructions in a robot interaction experiment. After gaining the attention of the person, in that experiment, the robot communicated to the person its requirement that the the person move off of the robot's path (i.e., to get out of the way). The use of this multimodal cue integrated with explicit instructions was found to be more effective than the robot just waiting for humans to move out of the way. In this earlier study we used a simple fan to provide the haptic air cues. Results presented here demonstrate that even more effective haptic air cues can be delivered through simple augmentation of the air source. If robots are to operate well in social settings, they require mechanisms to gain the attention of other people and to communicate their requirements to them. Haptic air cues provides a safe and socially acceptable mechanism for capturing that attention.

# REFERENCES


Burgoon, J. K., Johnson, M. L., and Koch, P. T. (1998). The nature and measurement of interpersonal dominance. *Communications Monographs*, 65(4):308–335.

Cattafesta, L., Bahr, C., and Mathew, J. (2010). *Fundamentals of Wind-Tunnel Design*. Wiley, Chichester, UK.

Dakopolos, A. J. and Jahromi, L. B. (2019). Differences in sensory responses among children with autism spectrum disorder and typical development: Links to joint attention and social competence. *Infant and Child Development*, 28(1):e2117. e2117 ICD-17-0123.R1.

Friedman, N., Goedicke, D., Jenkin, M., Degutyte, Z., Zhang, V., Astell, A., and Dudek, G. (2020a). Exploring modalities to ask, "excuse me, could i get through?". In *Proc. Workshop on Integrating Multidisciplinary Approaches to Advance Physical Human-Robot Interaction*, Paris, France. Held in conjunction with IEEE ICRA.

Friedman, N., Goedicke, D., Zhang, V., Jenkin, M., Degutyte, Z., Astell, A., Liu, S., and Dudek, G. (2020b). Out of my way! exploring different modalities for robots to ask people to move out of the way. In *Workshop on Active Vision and Perception in Human(-Robot) Collaboration*, Held Online. Held in conjunction with IEEE Int. Conf. on Robot and Human Interactive Communication.

Gupta, S., Morris, D., Patel, S., and Tan, D. (2013). AirWave: Non-contact haptic feedback using air vortex rings. In *ACM Joint Conference on Pervasive and Ubiquitous Computing*, pages 419–428, Zurich, Switzerland.

Hall, E. T. (1966). *The Hidden Dimension*, volume 609. Doubleday, Garden City, NY.

Hasegawa, K., Qiu, L., Noda, A., Inoue, S., and Shinoda, H. (2017). Electrically steerable ultrasound-driven long narrow air flow. *Applied Physics Letters*, 111:06104.

Hoshi, T., Iwamoto, T., and Shinoda, H. (2009). Non-contact tactile sensation synthesized by ultrasound transducers. In *3rd Joint Eurohaptics Conference and Symposium on Haptic Interfaces*, pages 256–260, Salt Lake City, UT.

Kant, K. and Zucker, S. (1986). Toward efficient trajectory planning: The path-velocity decomposition. *International Journal of Robotic Research - IJRR*, 5:72–89.

Long, B., Seah, S. A., Carter, T., and Subramanian, S. (2014). Rendering volumetric haptic shapes in mid-air using ultrasound. In *ACM SIGGRAPH Asia*, Shenzhen, China.

Morel, T. (1975). Comprehensive design of axisymmetric wind tunnel contractions. *Journal of Fluids Engineering, Trans. ASME*, 97(2):255.

Mowafi, O., Khamiodhi, M., and Abouelsaadat, W. (2015). Airdisplay: experimenting with air flow as a communication medium. In *INTERACT*, pages 316–323, Bamberg, Germany.

Sodhi, R., Poupyrev, I., Glisson, M., and Israr, A. (2013). AIREAL: Interactive tactile experiences in free air. In *ACM SIGGRAPH*, Anaheim, CA.

Suzuki, Y. and Kobayashi, M. (2005). Air jet driven force feedback in virtual reality. *IEEE Computer Graphics and Applications*, Jan/Feb:44–47.

Tolley, D., Nguyen, T., Tang, A., Ranasinghe, N., Kawauchi, K., and Yen, C.-C. (2019). Windywall: Exploring create wind simulation. In *TEI 2019: International Conference on Tangible and Embedded Interaction*, pages 635–644, Tempe, AZ.

Tsalamlal, M. Y., Issart, P., Ouarti, N., and Ammi, M. (2014). Hair: Haptic feedback with a mobile air jet. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2699–2706, Hong Kong, China. IEEE.

Verlinden, J., Mulder, F. A., Vergeest, J. S. M., de Jonge, B., Krutly, D., Nagy, Z. E., Logeman, B. J., and Schouten, P. (2013). Enhancement of presence in a virtual sailing environment through localized wind simulation. *Procedia Engineering*, 60:435–441.

Wogollter, M. S. (2006). Attention switch and maintenance. In Mahwah, N., editor, *Handbook of Warnings*, pages 245–265. Lawrence Erlbaum Associates.