nature machine intelligence

Supplementary information

https://doi.org/10.1038/s42256-024-00964-x

Causal chambers as a real-world physical testbed for AI methodology

In the format provided by the
authors and unedited

Supplementary Information

Appendix I

Additional Figures

Synthetic data

Ground truth + Gaussian noise

$\Delta P := \tilde{P}_{\mathrm{up}} - \tilde{P}_{\mathrm{dw}} = \beta_0 + \beta_1 \omega^2 + \epsilon$

with $\epsilon \sim \mathcal{N}(0, \sigma^2)$

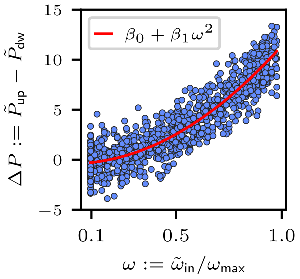

Estimates

5 random seeds

$$(R^2 = 0.71) \quad \Delta P = 14.0\omega - 3.9$$

$$(R^2 = 0.73) \quad \Delta P = 13.0\omega - 0.2 \tan\left(0.1 - \frac{8.9}{2.2\omega - 1.1}\right) - 3.4$$

$$(R^2 = 0.66) \quad \Delta P = 1.1 \arctan\left(-3.8 \tan\left(0.3 + \frac{58.2}{25.7\omega - 13.4}\right)\right) + 62.8\,|0.2\omega - 0.09|$$

$$(R^2 = 0.74) \quad \Delta P = (3.1\omega + 2.0)(3.2\omega - 1.7) - 4.2\left(0.1\left|35.6 \sin\left(\frac{0.3}{\cos(131.8\omega - 69.9)}\right) + 0.1\right| + 1\right)^{0.5} + 10.3$$

$$(R^2 = 0.76) \quad \Delta P = 7.3\omega + 14.3\,(\omega - 0.1 \sin(0.7 \tan(124.2\omega - 65.8) - 63.6) - 0.3)^2 - 1.8$$

Real data

Collected from the wind tunnel (same as Fig. 6e)

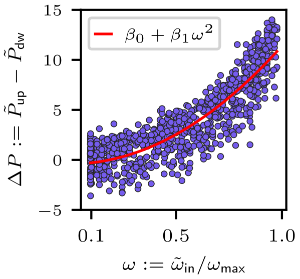

Estimates

5 random seeds

$$(R^2 = 0.75) \quad \Delta P = \exp\left(2.7\omega - \frac{0.8}{4.9 + \frac{1.4}{248622.0(\omega - 0.5)^4 - 5.8}}\right) + 0.1$$

$$(R^2 = 0.68) \quad \Delta P = (-2.3\omega - 0.2)(-5.8\omega + \cos(285.3\omega - 148.8) + 2.3) + 2.12$$

$$(R^2 = 0.75) \quad \Delta P = 10.1\omega^2 - 0.6015625\,(0.1 - \cos(81.0\,|340.4\omega - 171.0| + 4.5))^3 + 0.42$$

$$(R^2 = 0.72) \quad \Delta P = 6.7\omega^2 \left(1 + 0.4 \exp\left(-\frac{0.09}{4.2 - 3.1\,|31.6\omega - 16.5|}\right)\right)^2 + 1.2 \sin\left(4668785.1\,(\omega - 0.6)^2\right)$$

$$(R^2 = 0.72) \quad \Delta P = 7.2\left|(\omega + 0.2)^2 + 0.1 \sin(1.5\omega - 0.8)\right| - 0.09 \tan(14.0 \sin(19.6\omega - 10.3)) - 0.8$$

Supplementary Figure 1: Estimated expressions and their R^2 scores when we apply the symbolic regression method from Fig. 6e to recover Bernoulli's principle from synthetic data (top) and real data collected from the wind tunnel (bottom). The synthetic data is produced by fitting Bernoulli's principle to the data (see App. IV.1.3) and adding Gaussian noise to simulate sensor noise. For both datasets, the output varies with the random initialization of the method.

Appendix II

Chamber Variables

This appendix documents the variables in each chamber, that is, the actuators, sensor parameters, and sensor measurements. A description of the physical effects between them can be found in App. III.

II.1 Wind Tunnel

Variables	Value range	Type	Column name	Description
$L_{\text{in}}, L_{\text{out}}$	$[0, 1]$ (32-bit float)	A	`load_in` `load_out`	The load of the fans, corresponding to the duty cycle of the pulse-width-modulation (PWM) signal that controls their speed. For more details, see the datasheet for the `fan` component in App. VI.
$\tilde{C}_{\text{in}}, \tilde{C}_{\text{out}}$	$[0, 1023]$ (32-bit float)	S	`current_in` `current_out`	The uncalibrated measurements of the electric current drawn by the fans. The calibrated measurements (in amperes) are given by $\tilde{C}_{\text{in}} \times \frac{\text{vref}(R_{\text{in}})}{1023 \times 5} \times 2.5 \text{ and } \tilde{C}_{\text{out}} \times \frac{\text{vref}(R_{\text{out}})}{1023 \times 5} \times 2.5,$ where $\text{vref}(R_{\text{in}}), \text{vref}(R_{\text{out}})$ are the reference voltages of the corresponding sensors (see Table 3).
$\tilde{\omega}_{\text{in}}, \tilde{\omega}_{\text{out}}$	≥ 0 (32-bit float)	S	`rpm_in`, `rpm_out`	The speed of the fans in revolutions per minute.
$T_{\text{in}}, T_{\text{out}}$	$\{0,1\}$	P	`res_in`, `res_out`	The resolution of the tachometer timer that measures the elapsed time between successive revolutions of the fan, where 1 corresponds to microseconds and 0 to milliseconds. Choosing microseconds yields a higher resolution in the fan-speed measurement.

$\tilde{P}_{\text{up}}, \tilde{P}_{\text{dw}}$
$\tilde{P}_{\text{amb}}, \tilde{P}_{\text{int}}$
	≥ 0
	(32-bit float)
		S	`pressure_upwind`
			`pressure_downwind`
			`pressure_ambient`
			`pressure_intake`
				The barometric pressure, in pascals, as measured by the different barometers of the chamber. $\tilde{P}_{\text{int}}$ corresponds to the barometer placed at the tunnel intake. $\tilde{P}_{\text{amb}}$ is the ambient pressure measured by the outer barometer. $\tilde{P}_{\text{up}}$ and $\tilde{P}_{\text{dw}}$ correspond to the barometers inside the tunnel, which are placed facing into ($\tilde{P}_{\text{up}}$) and away ($\tilde{P}_{\text{dw}}$) from the airflow. Even at the same height and under the same conditions, the readings of the barometers have an offset resulting from the manufacturing process.

A_1, A_2	0:255	A	`pot_1`, `pot_2`	The wiper position of the digital potentiometers in the speaker amplification circuit (Fig. 2c).
$\tilde{S}_1, \tilde{S}_2$	$[0, 1023]$ (32-bit float)	S	`signal_1` `signal_2`	The amplitude of the signal after the first potentiometer ($\tilde{S}_1$) and after the second potentiometer ($\tilde{S}_2$) in the speaker amplification circuit (Fig. 2c). The calibrated amplitudes, in volts, are given by $\tilde{S}_1 \times \frac{\text{vref}(R_1)}{1023} \text{ and } \tilde{S}_2 \times \frac{\text{vref}(R_2)}{1023},$ where $\text{vref}(R_1), \text{vref}(R_2)$ are the reference voltages of the corresponding sensors (see Table 3).

Table 1: Description of the wind tunnel variables, including their symbol, value range, type, and the corresponding column name in the dataset files. Variables are categorized into different types: actuators (A) or sensor parameters (P), which can be directly manipulated, and sensor measurements (S), which cannot be directly manipulated. For the value range, $n{:}m$ corresponds to the range $\{n, n+1, \ldots, m\}$ and $n{:}m{:}k$ to $\{n, n+k, n+2k, \ldots, m\}$.

Variables	Value range	Type	Column name	Description
H	0:45:0.1	A	`hatch`	The position, in degrees, of the motor controlling the hatch that covers an additional opening of the wind tunnel. The position can be set in increments of 0.1° degrees, with the hatch being closed at 0° and fully open at 45°.
$\tilde{M}$	$[0, 1023]$ (32-bit float)	S	`mic`	The uncalibrated measurement of the sound level captured by the microphone. The calibrated signal amplitude, in volts, is given by $$\tilde{M} \times \frac{\text{vref}(R_M)}{1023},$$ where $\text{vref}(R_M)$ is the reference voltage of the corresponding sensor (see Table 3).

$R_{\text{in}}, R_{\text{out}}$
R_1, R_2
R_M
	$\{1.1, 2.56, 5\}$	P	`v_in`, `v_out`
			`v_1`, `v_2`
			`v_mic`
				The reference voltages, in volts, of the sensors used to measure the current $(\tilde{C}_{\text{in}}, \tilde{C}_{\text{out}})$, amplifier $(\tilde{S}_1, \tilde{S}_2)$ and microphone signals $(\tilde{M})$, respectively. The values differ slightly from the actual reference voltages seen by the sensors, which can be found in Table 3.

$O_{\text{in}}, O_{\text{out}}$
O_1, O_2, O_M
$O_{\text{up}}, O_{\text{dw}}$
$O_{\text{amb}}, O_{\text{int}}$
	{1,2,4,8}	P	`osr_in`, `osr_out`
			`osr_1`, `osr_2`
			`osr_mic`
			`osr_upwind`
			`osr_downwind`
			`osr_ambient`
			`osr_intake`
				The oversampling rates when taking measurements of the current $(\tilde{C}_{\text{in}}, \tilde{C}_{\text{out}})$, amplifier $(\tilde{S}_1, \tilde{S}_2)$ and microphone signals $(\tilde{M})$, and of air pressure at the different barometers $(\tilde{P}_{\text{up}}, \tilde{P}_{\text{dw}}, \tilde{P}_{\text{amb}}, \tilde{P}_{\text{int}})$. To avoid affecting the overall measurement time, the chambers always take the maximum number of readings (8) and discard excess readings accordingly.

Table 1: Description of the wind tunnel variables, including their symbol, value range, type, and the corresponding column name in the dataset files. Variables are categorized into different types: actuators (A) or sensor parameters (P), which can be directly manipulated, and sensor measurements (S), which cannot be directly manipulated. For the value range, n:m corresponds to the range $\{n, n+1, \ldots, m\}$ and n:m:k to $\{n, n+k, n+2k, \ldots, m\}$.

II.2 Light Tunnel

Variables	Value range	Type	Column name	Description

R, G, B	0:255	A	`red`
			`green`
			`blue`
				The brightness setting of the red, green, and blue LEDs on the main light source. Higher values correspond to higher brightness.

$\tilde{C}$	$[0, 1023]$ (32-bit float)	S	`current`	The uncalibrated measurement of the electric current drawn by the light source. The calibrated measurement (in amperes) is given by $$\tilde{C} \times \frac{\text{vref}(R_C)}{1023 \times 5} \times 2.5,$$ where $\text{vref}(R_C)$ is the reference voltage of the corresponding sensor (see Table 3).

$\tilde{I}_1, \tilde{I}_2, \tilde{I}_3$	$0{:}2^{16} - 1$	S	`ir_1`
			`ir_2`
			`ir_3`
				The uncalibrated infrared measurement of the light-intensity sensors, placed before $(\tilde{I}_1)$, between $(\tilde{I}_2)$, and after $(\tilde{I}_3)$ the polarizers, in reference to the light source.

Table 2: Description of the light tunnel variables, including their symbol, value range, type, and the corresponding column name in the dataset files. Variables are categorized into different types: actuators (A) or sensor parameters (P), which can be directly manipulated, and sensor measurements (S), which cannot be directly manipulated. For the value range, n:m corresponds to the range $\{n, n+1, \ldots, m\}$.

Variables	Value range	Type	Column name	Description

$\tilde{V}_1, \tilde{V}_2, \tilde{V}_3$	$0{:}2^{16}-1$	S	`vis_1`
			`vis_2`
			`vis_3`
				The uncalibrated visible-light measurement of the light-intensity sensors, placed before ($\tilde{V}_1$), between ($\tilde{V}_2$), and after ($\tilde{V}_3$) the polarizers, in reference to the light source.

D^I_1, D^I_2, D^I_3	$\{0,1,2\}$	P	`diode_ir_1`
			`diode_ir_2`
			`diode_ir_3`
				The photodiodes used by the light sensors to take infrared measurements, corresponding to the small ($D^I_j = 0$), medium ($D^I_j = 1$) and large ($D^I_j = 2$) infrared photodiodes onboard.

D^V_1, D^V_2, D^V_3	$\{0,1\}$	P	`diode_vis_1`
			`diode_vis_2`
			`diode_vis_3`
				The photodiodes used by the light sensors to take visible-light measurements, corresponding to the small ($D^V_j = 0$) and large ($D^V_j = 1$) photodiodes onboard.

T^I_1, T^I_2, T^I_3 T^V_1, T^V_2, T^V_3	$\{0,1,2,3\}$	P	`t_ir_1/2/3` `t_vis_1/2/3`	The exposure time of the photodiode during a light-intensity measurement.

L_{11}, L_{12},
L_{21}, L_{22},
L_{31}, L_{32}
	0:255	A	`l_11, l_12`
			`l_21, l_22`
			`l_31, l_32`
				The brightness of LEDs placed by each light-intensity sensor, where L_{i1}, L_{i2} correspond to the two LEDs placed by the i^{th} sensor ($\tilde{I}_i, \tilde{V}_i$). The value corresponds to the wiper position of a digital rheostat, which controls the current flowing to the LED. Higher values correspond to increased brightness.

θ_1, θ_2	-180:180:0.1	A	`pol_1` `pol_2`	The setting of the stepper motors that control the angle of the polarizer frames, which can be set in increments of 0.1 degrees. Because the mechanism functions without feedback, the actual angle of the polarizers may slightly deviate from this setting due to the imperfect coupling of the mechanical pieces.
$\tilde{\theta}_1, \tilde{\theta}_2$	$[0, 1023]$ (32-bit float)	S	`angle_1` `angle_2`	The position of the polarizers is encoded into a voltage using a rotary potentiometer, which is then read by the control computer to produce the measurements $\tilde{\theta}_1$ and $\tilde{\theta}_2$. Given these measurements, the calibrated angle measurement is given as $(\tilde{\theta}_j - Z_j) \times \frac{720}{1023} \times \frac{\text{vref}(R_j)}{5}$ degrees, where $\text{vref}(R_j)$ is the reference voltage of the corresponding sensor (see Table 3), and $Z_1 = 507, Z_2 = 512$ are the readings at angles $\theta_1 = \theta_2 = 0$ and reference voltages $R_1 = R_2 = 5$.

R_C, R_1, R_2	$\{1.1, 2.56, 5\}$	P	`v_c`
			`v_angle_1`
			`v_angle_2`
				The reference voltage, in volts, of the sensors used to measure the current ($\tilde{C}$) and polarizer angles ($\tilde{\theta}_1, \tilde{\theta}_2$), respectively. The values differ slightly from the actual reference voltages used by the sensors, which can be found in Table 3.

O_C, O_1, O_2	{1,2,4,8}	P	`osr_c`
			`osr_angle_1`
			`osr_angle_2`
				The oversampling rate of the sensors used to measure the current ($\tilde{C}$) and polarizer angles ($\tilde{\theta}_1, \tilde{\theta}_2$), respectively. To avoid affecting the overall measurement time, the chambers always take the maximum number of readings (8) and discard excess readings accordingly.

$\tilde{\text{Im}}$	24-bit RGB image	S	`im`	The color image produced by the camera, with a size of 2000×2000 pixels and 8 bits per color channel.
Ap	cf. Table 4	P	`aperture`	The f-number describing the aperture of the camera lens, with higher values corresponding to smaller openings.

Table 2: Description of the light tunnel variables, including their symbol, value range, type, and the corresponding column name in the dataset files. Variables are categorized into different types: actuators (A) or sensor parameters (P), which can be directly manipulated, and sensor measurements (S), which cannot be directly manipulated. For the value range, $n{:}m$ corresponds to the range $\{n, n+1, \ldots, m\}$.

Variables	Value range	Type	Column name	Description
ISO	cf. Table 4	P	`iso`	The gain of the camera sensor, where higher values correspond to higher sensitivity.
T_{Im}	cf. Table 4	P	`shutter_speed`	The shutter speed of the camera, i.e., how many seconds the camera sensor is exposed when taking an image.

Table 2: Description of the light tunnel variables, including their symbol, value range, type, and the corresponding column name in the dataset files. Variables are categorized into different types: actuators (A) or sensor parameters (P), which can be directly manipulated, and sensor measurements (S), which cannot be directly manipulated. For the value range, $n{:}m$ corresponds to the range $\{n, n+1, \dots, m\}$.

	Wind Tunnel	Light Tunnel
R	vref(R)	vref(R)
1.1	1.16	1.09
2.56	2.65	2.55
5	5	5

Table 3: Calibrated reference voltages for the wind tunnel ($R = R_{\mathrm{in}}, R_{\mathrm{out}}, R_1, R_2, R_M$) and the light tunnel ($R = R_C, R_1, R_2,$). They are estimated using the `analog_calibration` experiments in the `wt_test_v1` and `lt_test_v1` datasets.

Variable	Values
Ap	1.8, 2.0, 2.2, 2.5, 2.8, 3.2, 3.5, 4.0, 4.5, 5.0, 5.6, 6.3, 7.1, 8.0, 9.0, 10, 11, 13, 14, 16, 18, 20, 22
ISO	100, 125, 160, 200, 250, 320, 400, 500, 640, 800, 1000, 1250, 1600, 2000, 2500, 3200, 4000, 5000, 6400, 8000, 10000, 12800, 16000, 20000, 25600, 32000, 40000, 51200
T_{Im}	1/200, 1/250, 1/320, 1/400, 1/500, 1/640, 1/800, 1/1000, 1/1250, 1/1600, 1/2000, 1/2500, 1/3200, 1/4000

Table 4: Possible values for the camera parameters controlling the aperture of the lens (Ap), the sensor gain (ISO), and the shutter speed (T_{Im}).

Appendix III

Description of Physical Effects

We describe the effects that the chamber actuators and sensor parameters have on the measurements of each sensor. The effects correspond to the edges in the ground-truth graphs of the standard configurations in Fig. 3, where an edge $A \to B$ denotes that the actuator or sensor parameter A has an effect on the sensor measurement B. We justify each effect in terms of the chamber design and the underlying physical principles, further characterizing each effect through additional experiments (Supplementary Figure 2 to Supplementary Figure 12). As discussed in Sec. 3.1, these effects can be understood as causal effects—see App. V for an in-depth discussion and additional validation through randomized experiments. Throughout this appendix, we use a short-hand notation for denoting multiple effects e.g., $\{A_1, A_2\} \to \{B_1, B_2\}$ refers to the four edges $A_1 \to B_1, A_1 \to B_2, A_2 \to B_1$ and $A_2 \to B_2$.

The code to generate the plots in this section can be found in the Jupyter notebook `case_studies/plots_appendices.ipynb` in the paper repository `github.com/juangamella/causal-chamber-paper`.

III.1 Wind Tunnel

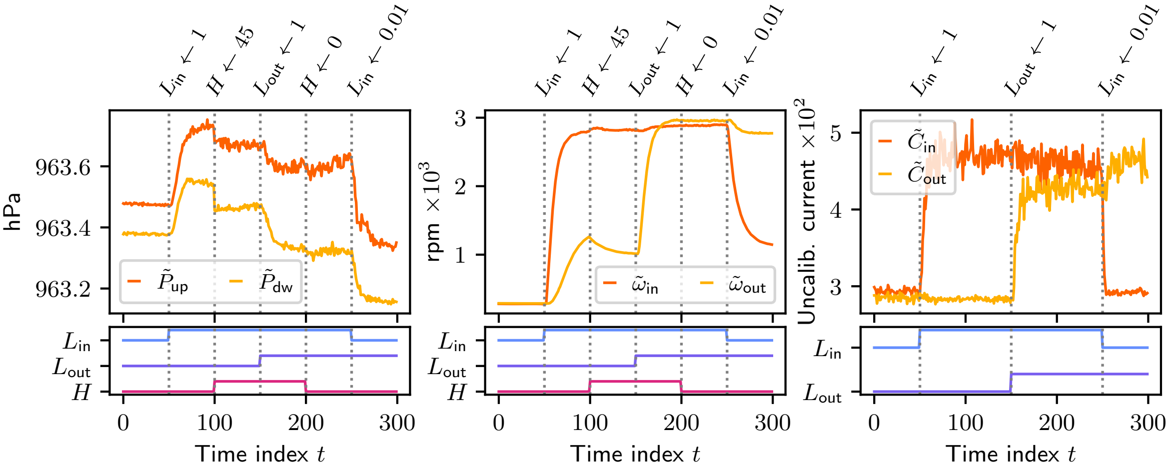

Supplementary Figure 2: Time-series data from different sensor measurements under varying inputs to the fan loads $(L_{\text{in}}, L_{\text{out}})$ and hatch position (H), shown on the bottom row. (left) The measurements $\tilde{P}_{\text{up}}, \tilde{P}_{\text{dw}}$ of the inner barometers are affected by both loads and the hatch position. (center) The fan speeds are affected by their corresponding load. Because the fans are placed in tandem, their speed is also affected by the load of the other fan, with the strength of the effect depending on the hatch position (e.g., $t = 100$). (right) Because they share the same power supply, the load of one fan has a slight effect on the current drawn by the other (e.g., $t = 250$). The data and corresponding experiment protocol can be found under the `steps` experiment in the `wt_test_v1` dataset.

$L_{\text{in}} \to \{\tilde{C}_{\text{in}}, \tilde{\omega}_{\text{in}}\}, L_{\text{out}} \to \{\tilde{C}_{\text{out}}, \tilde{\omega}_{\text{out}}\}$ The fan loads $(L_{\text{in}}, L_{\text{out}})$ define the duty cycle of the control signal sent to the fans, affecting their speed $(\tilde{\omega}_{\text{in}}, \tilde{\omega}_{\text{out}})$ and the current they consume $(\tilde{C}_{\text{in}}, \tilde{C}_{\text{out}})$. We provide mechanistic models describing these effects in App. IV.1.1, comparing their outputs to real measurements in Supplementary Figure 11. When the load is set to zero, the fan is completely powered off and no longer produces a tachometer signal (no corresponding observation shown in Supplementary Figure 11); the resulting speed measurement corresponds to the last measured speed (see Supplementary Figure 6).

$\{L_{\text{in}}, L_{\text{out}}, H\} \to \{\tilde{\omega}_{\text{in}}, \tilde{\omega}_{\text{out}}\}$ Each fan drives the flow of air through the wind tunnel, in turn making it easier or harder for the other fan to rotate. Thus, the speeds of the fans are coupled, and the strength of the coupling is affected by the hatch position (see Supplementary Figure 2, center).

$L_{\text{in}} \to \tilde{C}_{\text{out}}, L_{\text{out}} \to \tilde{C}_{\text{in}}$ Because they share the same power supply, the load of one fan has a small effect on the current drawn by the other (see Supplementary Figure 2, right). We believe these edges would not be present if each fan were driven by an independent power supply.

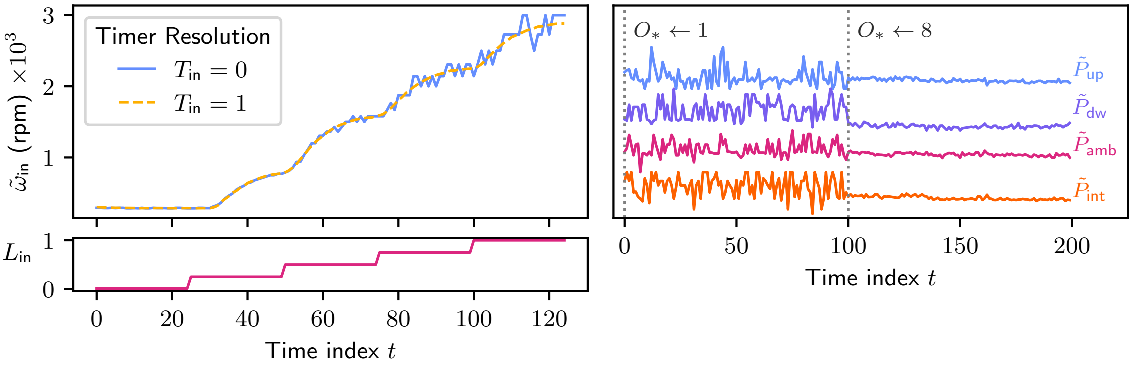

Supplementary Figure 3: (left) Measurements of the fan speed $\tilde{\omega}_{\text{in}}$ under a progressive increase in the fan load L_{in} for different resolutions of the tachometer timer, with ($T_{\text{in}} = 0$) corresponding to milliseconds and ($T_{\text{in}} = 1$) to microseconds. (right) Effect of the barometer oversampling rate ($O_{\text{up}}, O_{\text{dw}}, O_{\text{amb}}, O_{\text{int}}$) on the resulting measurement ($\tilde{P}_{\text{up}}, \tilde{P}_{\text{dw}}, \tilde{P}_{\text{amb}}, \tilde{P}_{\text{int}}$). For all barometers, the oversampling rate is increased from 1 to 8 at $t = 100$, while keeping all other chamber actuators and sensor parameters constant. The data and experimental setup for both plots can be found under the `tach_resolution` and `osr_barometers` experiments in the `wt_test_v1` dataset.

$T_{\text{in}} \rightarrow \tilde{\omega}_{\text{in}}, T_{\text{out}} \rightarrow \tilde{\omega}_{\text{out}}$ Changing the resolution ($T_{\text{in}}, T_{\text{out}}$) of the timers used in the fan tachometers also changes the resolution of the resulting speed measurement ($\tilde{\omega}_{\text{in}}, \tilde{\omega}_{\text{out}}$). Using a resolution of microseconds (e.g., $T_{\text{in}} = 1$) allows measuring smaller changes in the fan speed, with the difference being more noticeable at higher speeds (see Supplementary Figure 3, left).

$\{L_{\text{in}}, L_{\text{out}}, H\} \rightarrow \{\tilde{P}_{\text{up}}, \tilde{P}_{\text{dw}}\}$ By controlling the speed of the fans, the loads ($L_{\text{in}}, L_{\text{out}}$) affect the air pressure inside the wind tunnel, and thus the measurements of the two inner barometers ($\tilde{P}_{\text{up}}, \tilde{P}_{\text{dw}}$). By controlling the size of the additional opening to the outside, the hatch position (H) also affects the pressure inside the tunnel (Supplementary Figure 2). We provide mechanistic models describing these effects in App. IV.1.

$\{L_{\text{in}}, L_{\text{out}}, H\} \rightarrow \tilde{P}_{\text{int}}$ The fan loads and hatch position also affect the airflow and air pressure at the intake of the tunnel.

$O_{\text{up}} \rightarrow \tilde{P}_{\text{up}}, O_{\text{dw}} \rightarrow \tilde{P}_{\text{dw}}, O_{\text{amb}} \rightarrow \tilde{P}_{\text{amb}}, O_{\text{int}} \rightarrow \tilde{P}_{\text{int}}$ The oversampling rate determines how many barometric readings are averaged to produce a single measurement of the air pressure. A higher oversampling rate increases the precision of the barometers, reducing the noise in their measurements (see Supplementary Figure 3, right).

$\{R_{\text{out}}, O_{\text{out}}\} \rightarrow \tilde{C}_{\text{out}}, \{R_{\text{in}}, O_{\text{in}}\} \rightarrow \tilde{C}_{\text{in}}$ The current sensors encode their reading of the current as a voltage between 0 and 5 volts. This voltage is then read by the onboard computer, linearly mapping the range of $[0, R_{\text{in}}]$ ($[0, R_{\text{out}}]$) volts to $[0, 1023]$. Thus, reducing the reference voltages $R_{\text{in}}, R_{\text{out}}$ increases the resolution of the current measurements (see Supplementary Figure 7, right), but can also cause the measurements to saturate if the voltage surpasses the reference voltage. As for the barometers, the oversampling rate ($O_{\text{in}}, O_{\text{out}}$) determines the number of readings that are averaged to produce a single measurement, thus affecting the precision of the measurements (see Supplementary Figure 7, left).

$A_1 \rightarrow \{\tilde{S}_1, \tilde{S}_2\}, A_2 \rightarrow \tilde{S}_2$ The signal fed to the speaker amplification circuit (Fig. 2c) is binary white noise with a period of $40\mu s$: after every period the voltage is set to 0 or 5 volts with equal probability and is statistically independent of previous voltages. The two potentiometers act as controllable voltage dividers, affecting the amplitude of the signal at different points of the circuit. Thus, the first potentiometer (A_1) affects the measurements $\tilde{S}_1, \tilde{S}_2$, whereas the second potentiometer (A_2) only has an effect on the measurement $\tilde{S}_2$ (Supplementary Figure 4).

$\{R_1, O_1\} \rightarrow \tilde{S}_1, \{R_2, O_2\} \rightarrow \tilde{S}_2$ As for the current measurements, the reference voltages R_1, R_2 and oversampling rates O_1, O_2 have an effect on the resulting measurements $\tilde{S}_1, \tilde{S}_2$.

$A_1 \rightarrow \tilde{M}$ The setting of the first potentiometer (A_1) controls the amplitude of the signal that is sent to the speaker, affecting the sound-level measurement $\tilde{M}$ taken by the microphone (Supplementary Figure 4, left).

$\{L_{\text{in}}, L_{\text{out}}, H\} \to \tilde{M}$ The speed of the fans, controlled by the loads $L_{\text{in}}, L_{\text{out}}$, have an effect on the sound level $\tilde{M}$ measured by the microphone (see Supplementary Figure 5, left). The position of the hatch (H) affects the amount of air flowing through the exhaust and over the microphone, also affecting its reading (Supplementary Figure 5, center).

$\{R_M, O_M\} \to \tilde{M}$ As for the other sensors, the reference voltage affects the resolution of the microphone signal measurement and can induce saturation if the signal voltage rises above it. The oversampling rate O_M determines the Nyquist frequency of the system and can introduce aliasing in the resulting signal (Supplementary Figure 5, right).

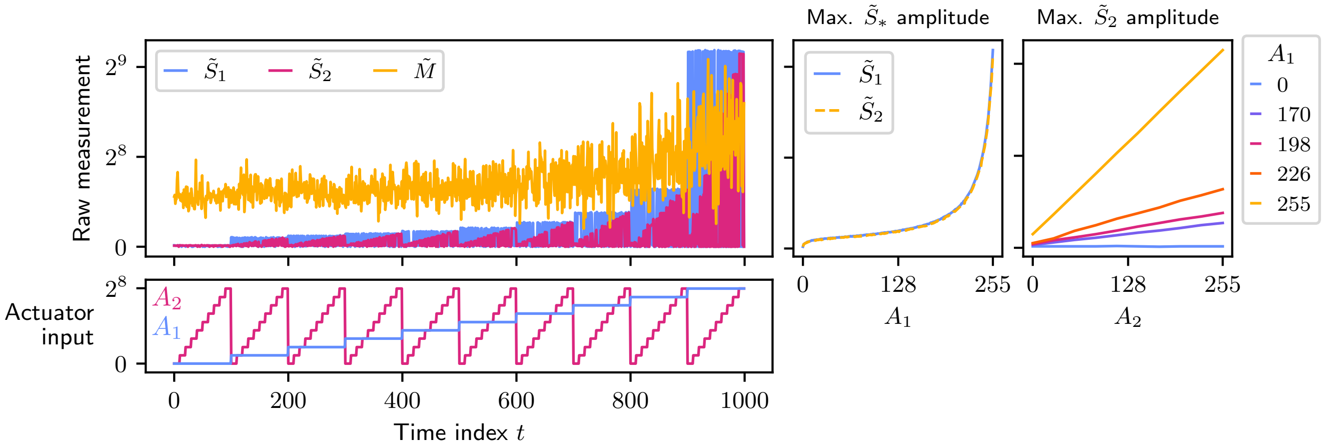

Supplementary Figure 4: (left) Time-series data of $\tilde{S}_1, \tilde{S}_2$ and the microphone output $\tilde{M}$, collected under the inputs A_1, A_2 to the amplification circuit potentiometers (bottom left), while all other actuators and sensor parameters are kept constant. The setting of the first potentiometer A_1 determines the maximum amplitude of the signals $\tilde{S}_1, \tilde{S}_2$ (center) and of the signal sent to the speaker, affecting the microphone output $\tilde{M}$ (left, gold). The setting of the second potentiometer A_2 only affects the maximum amplitude of $\tilde{S}_2$ (left), and the effect is linear for a fixed value of A_1 (right). The data and experimental setup for the plots can be found under the `potis_coarse` and `potis_fine` experiments in the `wt_test_v1` dataset.

$\tilde{P}_{\mathbf{dw}} \to \{\mathbf{L_{in}}, \mathbf{L_{out}}\}$ In the *pressure-control* configuration of the wind tunnel, the fan loads $L_{\text{in}}, L_{\text{out}}$ are set following a control mechanism to keep the pressure measured at the downwind barometer ($\tilde{P}_{\text{dw}}$) constant. In particular, the loads of the fans are set as

$$L_{\text{in}} \leftarrow \begin{cases} \min(1, u(t)) & \text{if } u(t) > 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad L_{\text{out}} \leftarrow \begin{cases} \min(1, -u(t)) & \text{if } u(t) < 0 \\ 0 & \text{otherwise} \end{cases}$$

where $u(t)$ is the controller output of the PID controller given by

$$u(t) := K_p e(t) + K_i \sum_{\tau=0}^{t} e(t) + K_d(e(t) - e(t-1)),$$

where $e(t) := T - \tilde{P}_{\text{dw}}^t$ is the control error, i.e., the difference in the pressure target T and the measurement $\tilde{P}_{\text{dw}}^t$ of the downwind barometer at time point t. For the `wt_pressure_control_v1` dataset, we set $(K_p, K_i, K_d) := (0.5, 0.1, 10^{-3})$ and set as pressure target T the first measurement $\tilde{P}_{\text{dw}}^0$ taken by the downwind barometer after the chamber powers up. The control mechanism is executed internally by the chamber computer, producing a new output $u(t)$ at every time step. The related variables, i.e., target T, control constants K_p, K_i, K_d, output $u(t)$ and error terms $e(t), \sum_{\tau=0}^{t} e(t)$ and $(e(t) - e(t-1))$, are returned with each measurement as additional columns in datasets collected from the pressure-control configuration.

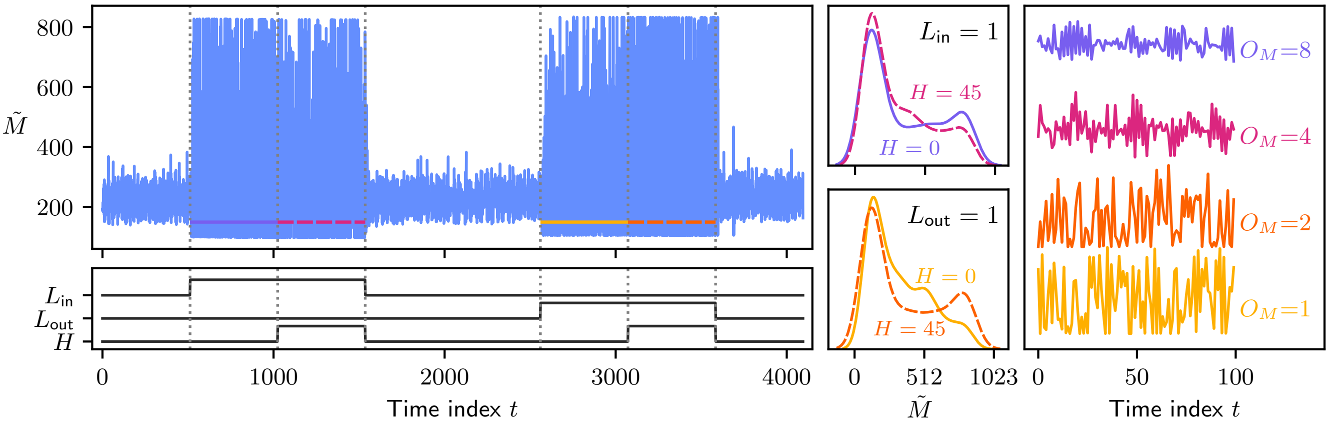

Supplementary Figure 5: (left) Time-series data of the microphone output $\tilde{M}$ collected under varying inputs to the fan loads $L_{\text{in}}, L_{\text{out}}$ and the hatch position H (bottom left). All three actuators have an effect on the microphone output. (center) Marginal distribution of the microphone output $\tilde{M}$ for $L_{\text{in}} = 1, L_{\text{out}} = 0$ (top) and $L_{\text{in}} = 0, L_{\text{out}} = 1$ (bottom), corresponding to the regions marked in the time-series plot on the left. The hatch affects the amount of air that flows through the exhaust of the tunnel and over the microphone, affecting its readings. The effect of opening the hatch ($H = 45$) is dependent on the loads of the fans, decreasing the airflow when $L_{\text{in}} = 1, L_{\text{out}} = 0$ and increasing it when $L_{\text{in}} = 0, L_{\text{out}} = 1$. (right) The oversampling rate determines the Nyquist frequency of the microphone system, introducing aliasing in the resulting signal. The data and experimental setup for the plots can be found in the `mic_effects` experiment in the `wt_test_v1` dataset.

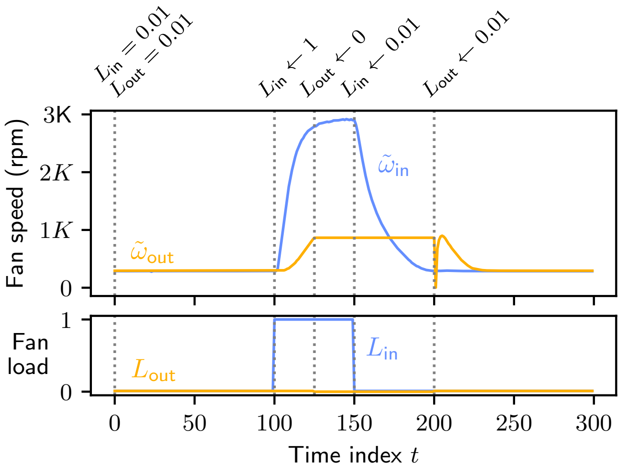

Supplementary Figure 6: By setting a fan load to zero (e.g., $L_{\text{out}} \leftarrow 0$ at $t = 125$), the fan is completely powered off and will decelerate until it stops rotating; however, at zero load, the fan no longer produces a tachometer signal, and the resulting speed measurement $\tilde{\omega}_{\text{out}}$ corresponds to the last measured speed (e.g., $t \in [125, 200]$). When powered up again ($t = 200$) the fan always draws full power for an instant, quickly accelerating before returning to the level specified by the load. The data and experimental setup for the plot can be found under the `zero_load` experiment in the `wt_test_v1` dataset.

III.2 Light Tunnel

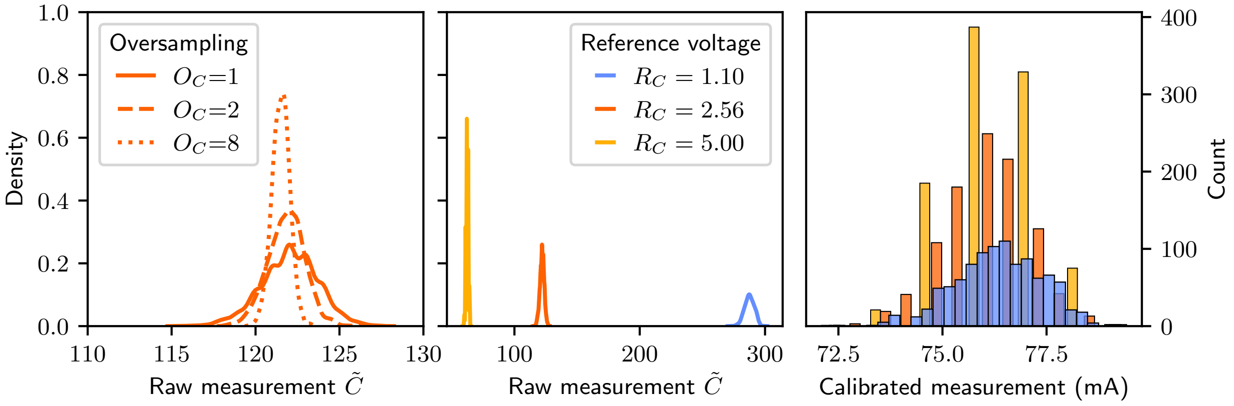

Supplementary Figure 7: Effect of the reference voltage R_C and oversampling rate O_C on the measurements $\tilde{C}$ of the current drawn by light source, when its brightness (R, G, B) is kept constant. The oversampling rate determines the amount of readings that are averaged to produce a single measurement, affecting the precision of the sensor (left). For a constant oversampling rate $(O_C = 1)$, changing the reference voltage results in a shift of the sensor output (middle) and an increase in the resolution of the calibrated measurement (right). While the plots show the current measurement $\tilde{C}$ in the light-tunnel, the same principles apply for the other sensors which encode their measurement as voltage $(\tilde{\theta}_1, \tilde{\theta}_2, \tilde{M}, \tilde{S}_1, \tilde{S}_2)$ and for the wind-tunnel barometers $(\tilde{P}_{\text{dw}}, \tilde{P}_{\text{up}}, \tilde{P}_{\text{amb}}, \tilde{P}_{\text{int}})$, which also allow for varying oversampling rates $(O_{\text{dw}}, O_{\text{up}}, O_{\text{amb}}, O_{\text{int}})$. The data corresponds to the `current_sensor` experiment from the `lt_test_v1` dataset.

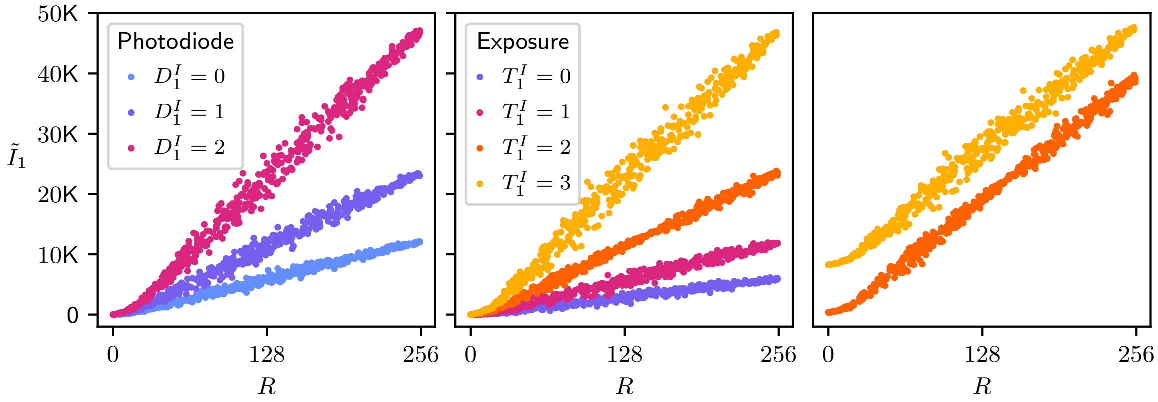

Supplementary Figure 8: Effect of the photodiode size (D^I_1) and exposure time (T^I_1) on the infrared measurements $(\tilde{I}_1)$ of the light sensor nearest the light source. We display the output $\tilde{I}_1$ of the sensor for varying brightness of the red LEDs on the light source (R), while keeping $G = B = 0$ and all other actuators constant. Increasing the photodiode size (left) or exposure time (center) both increase the sensitivity of the sensor. Due to an interplay between the exposure time and the pulse-width-modulation frequency of the light source, changes in the exposure time also affect the conditional distribution of the measurements given the light-source intensity (R, G, B). For example, when comparing the measurements for $T^I_1 = 2$ and $T^I_1 = 3$ in the rightmost panel (same data as in the center panel, but vertically scaled for comparison), we can see a difference in the variance as a function of R. The effects on the visible-light measurements $\tilde{V}_1$ follow the same principles. The data corresponds to the `ir_sensors` experiment from the `lt_test_v1` dataset.

$\{R, G, B\} \rightarrow \{\tilde{I}_1, \tilde{I}_2, \tilde{I}_3, \tilde{V}_1, \tilde{V}_2, \tilde{V}_3\}$ The brightness settings of the light-source colors (R, G, B) affect the readings of all light-intensity sensors. The effect of each color is approximately linear with heteroscedastic noise (Fig. 4d), with the slope determined by its typical wavelength and the spectral sensitivity of the sensor (see `light_source` and `light_sensor` components in App. VI).

$\{R, G, B\} \rightarrow \tilde{C}$ The brightness settings also affect the electric current drawn by the light source. The effect is again approximately linear with roughly the same slope for each channel (see Fig. 4d).

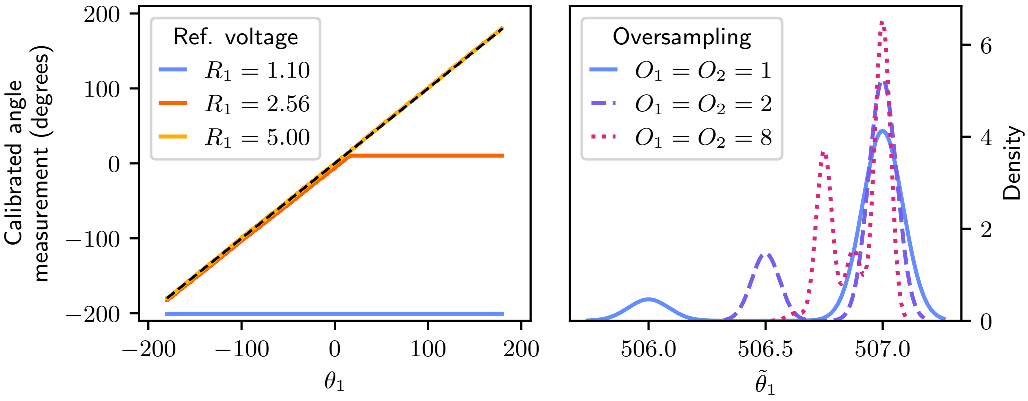

Supplementary Figure 9: Effect of the reference voltage (R_1) and oversampling rate (O_1) on the angle measurements of the first polarizer ($\tilde{\theta}_1$). (left) We show the calibrated measurements of the angle sensor (see Table 2 for details) for different values of the motor setting θ_1 and reference voltages R_1, while keeping the oversampling rate constant at $O_1 = 1$. Lowering the reference voltage R_1 results in a saturation of the sensor, being completely saturated (i.e., returning a constant reading) at $R_1 = 1.1$. (right) Uncalibrated measurements $\tilde{\theta}_1$ for a constant motor setting $\theta_1 = 0$ and reference voltage $R_1 = 5$. The stochasticity of the measurement is due to electrical noise in the voltage measurement—increasing the oversampling rate O_1 increases the precision of the sensor. The same principles apply to the measurements $\tilde{\theta}_2$ of the second angle sensor. The data corresponds to the `angle_sensors` experiment from the `lt_test_v1` dataset.

$\{R_C, O_C\} \rightarrow \tilde{C}$ As for the wind-tunnel, the sensor measuring the current drawn by the light source encodes its reading into a voltage between 0 and 5 volts. This voltage is then read by the onboard computer, linearly mapping the range of $[0, R_C]$ volts to $[0, 1023]$. Thus, reducing the reference voltage R_C increases the resolution of the current measurements (see Supplementary Figure 7, right). The oversampling rate (O_C) determines the number of readings that are averaged to produce a single measurement, affecting its precision (Supplementary Figure 7, left).

$\{D^I_j, T^I_j\} \rightarrow \tilde{I}_j, \{D^V_j, T^V_j\} \rightarrow \tilde{V}_j$ The measurements of each light sensor are affected by the choice of photodiode (D^I_j, D^V_j) and gain (T^I_j, T^V_j) used to perform the infrared ($\tilde{I}_j$) and visible-light ($\tilde{V}_j$) readings. Larger photodiodes collect light over a larger area, increasing the sensitivity of the sensor (Supplementary Figure 8, left). Higher exposure times achieve a similar effect by increasing the reading duration, which also has an effect on the conditional distribution of the measurements given the light source brightness (Supplementary Figure 8, right). A timing mechanism ensures that changes in the exposure time do not affect the overall measurement time.

$\{\theta_j, R_j, O_j\} \rightarrow \tilde{\theta}_j$ The motor settings θ_1, θ_2 determine the positions of the polarizer frames. The positions are encoded into a voltage by means of rotary potentiometers, resulting in the measurements $\tilde{\theta}_1, \tilde{\theta}_2$. Changes in the reference voltages R_1, R_2 affect the resolution of the measurements, but can also cause the sensor to saturate (see Supplementary Figure 9, left). The oversampling rates (O_1, O_2) determine the number of readings that are averaged to produce a single measurement, affecting its precision (Supplementary Figure 9, right).

$\{\theta_1, \theta_2\} \rightarrow \{\tilde{I}_3, \tilde{V}_3, \tilde{\mathrm{Im}}\}$ The position of the polarizers affects the intensity and spectral composition (i.e., color) of the light passing through them, affecting the readings ($\tilde{I}_3, \tilde{V}_3$) of the third light sensor (see Fig. 4c and Supplementary Figure 12) and the images ($\tilde{\mathrm{Im}}$) captured by the camera (see Fig. 4e). The effect is described by Malus' law—see App. IV.2.1 for more details.

$\{\mathrm{Ap}, \mathrm{ISO}, T\} \rightarrow \tilde{\mathrm{Im}}$ The camera parameters affect the brightness of the captured images while introducing a variety of side effects (Fig. 4e). For example, a higher gain (ISO) increases the noise, while the aperture (Ap) changes the depth of field, introducing blur.

$\{L_{j1}, L_{j1}\} \rightarrow \tilde{I}_j, \tilde{V}_j$ Besides the light source, two additional LEDs placed by each light sensor (see Fig. 2d) have an effect on its reading. To avoid affecting the measurements of the other light sensors, the LEDs only turn on when their corresponding sensor is taking a measurement. Their brightness is controlled by means of digital rheostats that control the current flowing to each LED. The settings L_{j1}, L_{j2} correspond to the wiper position of each rheostat. The relationship between the setting and the intensity readings follows an exponential function (Fig. 4b), resulting from the LED's voltage-current characteristic, luminosity response and typical wavelength, and the spectral sensitivity of the sensor. The brightness of the LEDs is much lower than that of the light source (Supplementary Figure 10). See the `led` and `light_sensor` components in App. VI for the corresponding datasheets.

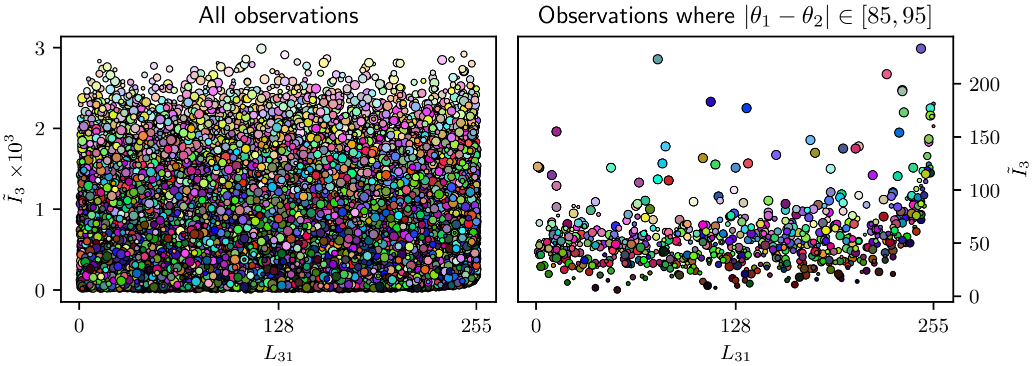

Supplementary Figure 10: Effect of the LED brightness settings (L_{31}, L_{32}) on the reading $\tilde{I}_3$ of the third light sensor, which is separated from the light source by both polarizers. The light source colors R, G, B and brightness settings $L_{11}, \ldots L_{32}$ are sampled independently and uniformly at random from the range $\{0, \ldots, 255\}$, and the polarizer angles θ_1, θ_2 from $[-15, 105]$. In both plots, the color corresponds to R, G, B, and the size to L_{32}. Because the brightness of the LEDs is very small when compared to that of the light source, their effect on the intensity reading $\tilde{I}_3$ is barely noticeable when looking at all observations (left), but becomes apparent, for example, when the polarizers are close to orthogonal (right) and block most of the light emanating from the light source. The data corresponds to the `actuators_white` experiment from the `lt_walks_v1` dataset.

Appendix IV

Mechanistic Models of The Chambers

This appendix provides mechanistic models describing several effects and processes in the causal chambers. Each model is labeled with a letter and a number (e.g., B1, C2); different letters correspond to different modeled quantities, and higher numbers mean increased fidelity. A Python implementation of each model can be found in `github.com/juangamella/causal-chamber`, and is accessible through the `causalchamber` package; examples and the code to generate the plots in this section can be found in the Jupyter notebook `case_studies/mechanistic_models.ipynb` in the paper repository `github.com/juangamella/causal-chamber-paper`.

Excluding models D1 and F1–F3, the models in this section are agnostic to phenomena arising from the sensors that measure the modeled physical quantities, such as measurement and quantization noise, saturation, or non-linearities in their response. For this reason, we denote the modeled quantities without a tilde (e.g., $\omega_{\text{in}}, P_{\text{dw}}, I_3$) to differentiate them from the corresponding sensor measurements (e.g., $\tilde{\omega}_{\text{in}}, \tilde{P}_{\text{dw}}, \tilde{I}_3$).

IV.1 Wind Tunnel

We provide models describing the effect of the fan load on the speed of the fans and the drawn current (App. IV.1.1); the effect of the fan loads and hatch position on the reading of the downwind barometer (App. IV.1.2); and the difference between the readings of the up- and downwind barometers, as described by Bernoulli's principle (App. IV.1.3).

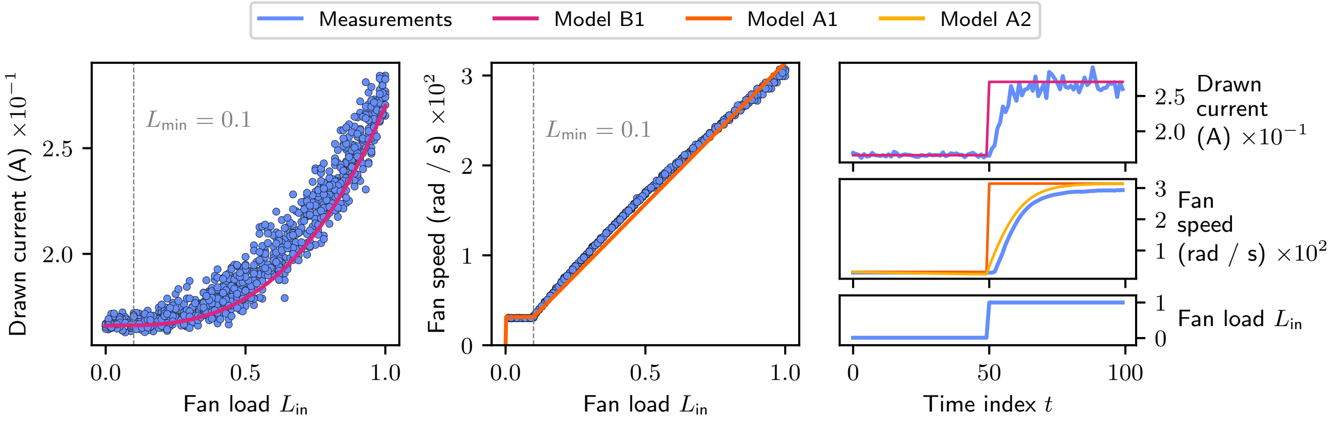

Supplementary Figure 11: Comparing model predictions with real measurements. (left and center): models B1 and A1 of the drawn current and speed of the fans, respectively, as a function of the load L_{in}, compared to steady-state measurements from the `wt_bernoulli_v1` dataset. No data is shown for $L_{\text{in}} = 0$ as the fan is completely powered off and no longer produces a tachometer signal—see Supplementary Figure 6. (right) Output of the models B1, A1 and A2 for time-series data from the `wt_test_v1` dataset. The parameters of all models are set as in App. IV.1.1.

The fans used in the wind tunnel are high-speed fans for industrial cooling applications. Their speed is controlled via pulse-width-modulation (PWM), where the pulse-width, or duty cycle, corresponds to the fan loads L_{in} and L_{out} of the wind tunnel. Additional details and technical specifications of the fans can be found in the datasheets of App. VI, under the `fan` component. As we derive the different models of this section, we will refer to these datasheets for the values of certain model parameters.

IV.1.1 Effect of the fan load on fan speed and drawn current

Two important aspects of the fans' design are used throughout the models of the fan speed (A1, A2) and drawn current (B1): first, the fans are designed so their steady-state speed scales broadly linearly with the load. At full load, the fans turn at a maximum speed ω_{max}, which is specified by the manufacturer as 314.16 rad/s (3000 RPM). Second, due to their intended application, unless completely powered off (i.e., $L_{\text{in}}, L_{\text{out}} = 0$) the fans never operate below a certain speed, corresponding to a minimum effective load which we denote by L_{min}. This value is specified by the manufacturer to be around 0.2, but our experiments show it to be closer to 0.1 (see Supplementary Figure 11, center).

Model A1

We model the steady-state speed ω of the fan under a load L as

$$\omega = \begin{cases} \max(L, L_{\min})\omega_{\max} & \text{if } L > 0 \\ 0 & \text{if } L = 0 \end{cases} \tag{2}$$

where $\Theta := (\omega_{\max}, L_{\min})$ are the parameters of the model. The first two correspond to the maximum speed of the fan and the minimum effective load. From the datasheets and our experiments, we can set $L_{\min} = 0.1$ and $\omega_{\max} = 314.16$ rad/s (see Supplementary Figure 11, left).

Model B1

We model the effect of fan load on its drawn current through the affinity laws [e.g., 1, Sec. 9.3.5], which state that the power W consumed by a fan is proportional to the cube of its speed ω, that is, $W = W_{\max}(\omega/\omega_{\max})^3$, where $W_{\max}$ is the power consumed by the fan at its maximum speed $\omega_{\max}$. Because the fans operate at a constant voltage, we can rewrite the above expression in terms of the drawn current C as $C = C_{\max}(\omega/\omega_{\max})^3$, where $C_{\max}$ is the current drawn at maximum speed. Accounting for the no-load current $C_{\min}$ of the motor, and combining the above with model A1 of the fan speed, we arrive at the model

$$C = \begin{cases} C_{\min} + \max(L_{\min}, L)^3(C_{\max} - C_{\min}) & \text{if } L > 0 \\ C_{\min} & \text{if } L = 0 \end{cases} \tag{3}$$

where $\Theta := (L_{\min}, C_{\min}, C_{\max})$ are the parameters of the model. As before, we can set $L_{\min} = 0.1$; from the datasheet, we can set the nominal current of the fan to $C_{\max} = 0.26$ A, and empirically determine the no-load current to be approximately $C_{\min} = 0.166$ A.

Model A2

To model the dynamics of the fan speed beyond the steady state, we express the change in the speed through the torque-balance equation

$$d\omega = \frac{1}{I}(\tau(L) - K\omega^2)dt, \tag{4}$$

where I is the fan's moment of inertia, $\tau(L)$ is the torque applied by the fan motor for a given load L, and $K\omega^2$ is the opposing torque due to the drag experienced by the fan. Here, K is a constant that depends on the geometry of the blades and the density of the air [e.g., 1, Sec. 9.3.5]. These constitute the parameters $\Theta := (I, \tau(L), K)$ of our model, which we approximate as follows:

- We model the fan as a solid cylinder, resulting in a moment of inertia $I = \frac{1}{2}mr^2$, where m, r are the mass and radius of the disc, respectively. The fan has a radius of $r = 0.059$m, and we estimate its mass as $m = 0.02$kg, resulting in $I = 3.48 \times 10^{-5}\text{kg/m}^2$.
- We approximate the load-to-torque function as

  $$\tau(L) := \begin{cases} T(\max(L_{\min}, L)^3(C_{\max} - C_{\min})) & \text{if } L > 0 \\ 0 & \text{if } L = 0 \end{cases}$$

  where T is the torque constant of the motor and the remaining terms correspond to the drawn current as described by model B1 (after substracting the no-load current $C_{\min}$). In Supplementary Figure 11 and Fig. 6f, we set $T = 0.05$Nm/A.
- To obtain a value for the constant K in our model, we solve the steady-state equation $d\omega = \frac{1}{I}(\tau(1) - K\omega_{\max}^2) = 0$ with I and $\tau(L)$ as above, resulting in $K = 5.26 \times 10^{-8}$.

In Supplementary Figure 11 we compare the output of the models A1 and A2 to real measurements collected from the chamber. An important source of misspecification is that we model the intake and exhaust fans independently, whereas in the tunnel, their speeds affect each other (see Fig. 6f).

IV.1.2 Effect of the fan loads and hatch position on the downwind barometer

In this section, we provide models describing the effect of the fan loads and hatch position on the static pressure inside the wind tunnel, corresponding to the measurement $\tilde{P}_{\text{dw}}$ of the downwind barometer. Models C1–C3 relate the pressure to the fan speeds ω_{in} and hatch position H. We can combine them with models A1 or A2 of the fan speeds to simulate the barometer reading as a function of the fan loads $(L_{\text{in}}, L_{\text{out}})$, as we do in Fig. 6f.

To model the complex dynamics of the airflow through the wind tunnel, models C1–C3 make some simplifying assumptions. In first place, we assume that the change in static pressure can be simply computed as the difference of the static pressure produced by each fan. Furthermore, we model each fan independently, ignoring mutual effects between them.

Model C1

As a first approach, we model the static pressure P_{dw} inside the tunnel as

$$P_{\text{dw}} = P_{\text{amb}} + S_{\text{max}} \left(\frac{\omega_{\text{in}}}{\omega_{\text{max}}} \right)^2 - S_{\text{max}} \left(\frac{\omega_{\text{out}}}{\omega_{\text{max}}} \right)^2, \tag{5}$$

where P_{amb} is the static pressure outside the wind tunnel (corresponding to the barometer reading $\tilde{P}_{\text{amb}}$). The other two terms on the RHS correspond to the static pressures produced by the intake and exhaust fans, resulting from the the affinity laws [e.g., 2, chapter 5.14]. The parameters of our model are $\Theta := (S_{\text{max}}, \omega_{\text{max}})$, corresponding to the maximum static pressure and maximum speed of the fan. From the technical specifications, we can set these values to 74.82 Pa and 314.16 rad/s, respectively.

Model C2

Model C1 assumes that the fans always produce their maximum static pressure $S_{\text{max}} \left(\frac{\omega}{\omega_{\text{max}}} \right)^2$ at a given speed ω. In reality, there is a trade-off between the airflow and static pressure produced by the fan, dictated by the system's resistance to the flow of air, which is called *impedance*. As an illustration, a fan blowing into an enclosure with no other openings produces no airflow but maximum static pressure, whereas a fan at the boundary of two completely open enclosures produces its maximum airflow but no static pressure [3]. To account for this effect, we model the pressure P as

$$P_{\text{dw}} = P_{\text{amb}} + S_Z(\omega_{\text{in}}) - S_Z(\omega_{\text{out}}),$$

where $S_Z(\omega)$ is the static pressure produced by the fan at speed ω in a system with impedance Z. We compute $S_Z(\omega)$ as the intersection of the impedance curve $S = ZQ^2$ with the pressure-airflow characteristic (PQ-curve) of the fan at speed ω [see, e.g., 3]. Because the PQ-curve of our fans is not available, we approximate them as the linear relation

$$S = \left(\frac{\omega_{\text{in}}}{\omega_{\text{max}}} \right)^2 S_{\text{max}} - \left(\frac{\omega_{\text{in}}}{\omega_{\text{max}}} \right) \frac{S_{\text{max}}}{Q_{\text{max}}} Q,$$

where Q_{max} is the maximum airflow of the fan. The impedance $Z \in [0, \infty]$ is a result of the complex dynamics of the air as it flows through the wind tunnel, and we cannot directly calculate it. For a more intuitive interpretation, we can express it in terms of the ratio r of the maximum airflow that the fans produce when turning at full speed, yielding $Z := \frac{S_{\text{max}}}{Q_{\text{max}}^2} \left(\frac{1-r}{r^2} \right)$ for $r \in [0, 1]$. Under this parametrization, our model becomes

$$P_{\text{dw}} = P_{\text{amb}} + S_r(\omega_{\text{in}}) - S_r(\omega_{\text{out}}), \tag{6}$$

with parameters $\Theta := (S_{\text{max}}, \omega_{\text{max}}, Q_{\text{max}}, r)$. Following the fan's specifications, we set the maximum airflow to $Q_{\text{max}} = 0.052\text{m}^3/\text{s}$. For the simulation results in Fig. 6f, we set $r = 0.7$.

Model C3

So far. we have not considered the effect of the hatch position (H) on the pressure inside the tunnel. Building on Model C2, we model the effect of the hatch as a change in the impedance of the system. Opening the hatch lowers the system impedance and, thus, the static pressure produced by the fans. As in (6), we model the pressure as

$$P_{\text{dw}} = P_{\text{amb}} + S_{r(H)}(\omega_{\text{in}}) - S_{r(H)}(\omega_{\text{out}}), \tag{7}$$

except now the maximum airflow ratio is given by the function $r(H) = \min(1, r_0 + \beta \frac{H}{45})$, where $H \in [0, 45]$, $r_0 \in [0, 1]$ is the ratio when the hatch is closed $(H = 0)$, and $\beta > 0$ is the effect of the hatch, which we model as linear. Thus, the parameters of the model become $\Theta := (S_{\text{max}}, \omega_{\text{max}}, Q_{\text{max}}, r_0, \beta)$. For the simulation results in Fig. 6f, we set $r_0 = 0.7$, $\beta = 0.15$, and all other parameters as for the previous models.

IV.1.3 Relationship between the up- and downwind barometers

Due to their positioning, the pair of barometers inside the wind tunnel function as a pitot tube, where the upwind barometer ($\tilde{P}_{\text{up}}$) measures the stagnation or total pressure, and the downwind barometer ($\tilde{P}_{\text{dw}}$) measures the static pressure. Ignoring measurement noise, the readings of both barometers can be related by Bernoulli's principle:

$$\tilde{P}_{\text{up}} = \tilde{P}_{\text{dw}} + \frac{1}{2}\rho v^2 + \Delta. \tag{8}$$

Above, ρ is the density of the air, v is the speed of the airflow and Δ is the offset between the readings of the barometers due to manufacturing differences, which we empirically determine to be around $\Delta = 7.1$ Pa. By solving for v, we can use (8) to estimate the speed of the airflow through the tunnel. Model D1 relates the difference in the barometer readings to the speed of the intake fan.

Model D1

To express the speed v of the airflow as a function of the intake fan speed, we again rely on the affinity laws, which state that the airflow Q produced by the fan is proportional to its speed ω_{in}, i.e.

$$Q = \left(\frac{\omega_{\text{in}}}{\omega_{\text{max}}}\right) Q_{\text{max}},$$

where $\omega_{\text{max}}, Q_{\text{max}}$ are the maximum speed and airflow of the fan. Dividing by the area A of the fan opening, we obtain the following expression for airspeed

$$v = \frac{1}{A}\left(\frac{\omega_{\text{in}}}{\omega_{\text{max}}}\right) Q_{\text{max}}.$$

Combining this with (8), we arrive at

$$\tilde{P}_{\text{up}} - \tilde{P}_{\text{dw}} = \frac{\rho}{2A^2}\left(\frac{Q_{\text{max}}}{\omega_{\text{max}}}\right)^2 \omega_{\text{in}}^2 + \Delta, \tag{9}$$

where $\Theta := (\rho, A, \omega_{\text{max}}, Q_{\text{max}}, \Delta)$ are the parameters of our model. Given the specifications of the fan, we have $A = \pi 0.06^2$ m^2 and $\omega_{\text{max}} = 314.16$ rad/s. We can set $\rho = 1.2$ and we empirically determine $\Delta = 7.1$ Pa. The same difficulties apply to estimating Q_{max} as for S_{max} in model C1. Furthermore, in this model, we ignore the additional effect that the exhaust fan has on the total airflow through the tunnel. For the symbolic regression task in Fig. 6e, we turn off the exhaust fan and take steady-state measurements for different loads of the intake fan (see the `wt_bernoulli_v1` dataset for more details).

IV.2 Light Tunnel

In App. IV.2.1, we provide some background on Malus' law and provide a model of the effect of the polarizer angles (θ_1, θ_2) on the light intensity reaching the third light sensor ($\tilde{I}_3, \tilde{V}_3$). In App. IV.2.2 we provide models of the image-capture process with increasing degrees of fidelity.

IV.2.1 Polarizer Effects

Depending on their angle, the linear polarizers dim the light passing through them. Malus' law describes their effect on light intensity [4]: given a beam of totally polarized light and a perfect polarizer, the intensity I after the polarizer is given by $I = I_0 \cos^2(\theta)$, where I_0 is the intensity before the polarizer and θ is the angle between the polarizer axis and the polarization of the incident light. For unpolarized light passing through two polarizers at angles θ_1 and θ_2, the intensity after the polarizers becomes

$$I = \frac{1}{2} I_0 \cos^2(\theta_1 - \theta_2).$$

In practice, polarizers are imperfect: they are not fully transparent to light polarized parallel to their polarization axis, and they do not block all light with polarization orthogonal to it. Furthermore, their effect depends on the frequency of light. For an unpolarized light beam with a given spectral composition, a pair of polarizers will scale its intensity by $T^p < 1/2$; when orthogonal, they will still allow some light to pass through, scaling the intensity by $T^c > 0$.

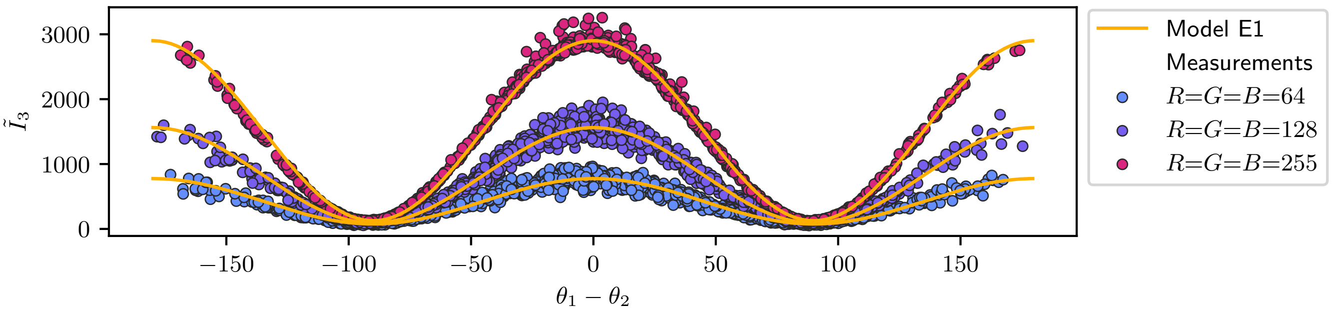

Supplementary Figure 12: Comparing the prediction of model E1 (Malus' law) with real measurements gathered from the light tunnel. Model E1 is fit to data gathered under different brightness settings of the light source (i.e., $R = G = B = 64, 128, 255$), obtaining a coefficient of determination of $R^2 = 0.93, 0.97, 0.99$, respectively. The data for $R = G = B = 255$ is used in the symbolic regression task of Fig. 6e. The signal-to-noise decreases with the light-source brightness, allowing for more challenging scenarios. The data can be found in the `lt_malus_v1` dataset.

Model E1

To model the effect of the polarizer angles θ_1, θ_2 on the intensity I_3 of light reaching the third sensor, we can reformulate the generalization of Malus' law for imperfect polarizers [5] in terms of the transmission rates T^p, T^c, resulting in the model

$$I_3 = I_0[(T^p - T^c)\cos^2(\theta_1 - \theta_2) + T^c],$$

where $\Theta := (I_0, T^p, T^c)$ are the parameters of the model. The radiant intensity before the polarizers (I_0) is given in watts per steradian (W/sr), and depends e.g., on the brightness of the tunnel light source. For the polarizing film used in the light tunnel, values for the transmission rates $T^p, T^c \in [0, 1]$ can be found in App. VI under the `polarizer` component. We can rewrite the above model as

$$I_3 = \beta_1 \cos^2(\theta_1 - \theta_2) + \beta_0,$$

becoming our ground truth for the symbolic regression task in Fig. 6e. We compare, in Supplementary Figure 12, the output of model E1 to measurements gathered from the chamber, with the parameters β_0, β_1 fit to the data.

IV.2.2 Image Capture

As models of the image-capture process in the light tunnel, we provide simple simulators that produce a synthetic image given the light-source setting R, G, B and polarizer angles θ_1, θ_2. The resulting image simulates the image ($\tilde{\mathrm{Im}}$) produced by the light tunnel and consists of a hexagon over a black background, whose color is given by the RGB vector $(\tilde{R}, \tilde{G}, \tilde{B}) \in [0, 1]^3$ and models the response of the camera sensor. The models F1–F3 differ in how the vector $(\tilde{R}, \tilde{G}, \tilde{B})$ is computed. In Fig. 6f, we compare their output to real images produced by the light tunnel.

Model F1

The first model assumes the linear polarizers to be perfect and that their effect is uniform across all wavelengths. Furthermore, it assumes that the camera sensor is perfectly calibrated to the light source. The resulting color vector is given by

$$\begin{pmatrix} \tilde{R} \\ \tilde{G} \\ \tilde{B} \end{pmatrix} = \cos^2(\theta_1 - \theta_2) \frac{1}{255} \begin{pmatrix} R \\ G \\ B \end{pmatrix}, \tag{10}$$

where $\cos^2(\theta_1 - \theta_2)$ models the dimming effect of the polarizers according to Malus' law (see App. IV.2.1 above).

Model F2

The second model explicitly models the output of the camera sensor as a result of its spectral sensitivity and the white balance correction applied by the processor in the camera. Our complete model is

$$\begin{pmatrix} \tilde{R} \\ \tilde{G} \\ \tilde{B} \end{pmatrix} = \min\{1, eWS\cos^2(\theta_1 - \theta_2)\frac{1}{255}\begin{pmatrix} R \\ G \\ B \end{pmatrix}\}, \tag{11}$$

where

- $S \in \mathbb{R}_+^{3\times 3}$ models the response of the sensor to the light of each color,
- $W := \operatorname{diag}(w_R, w_B, w_C)$ with $w_R, w_B, w_C \in \mathbb{R}_+ : w_R + w_B + w_C = 1$ is the white-balance correction applied by the camera, and
- $e > 0$ is an additional parameter to model the exposure of the sensor as affected by the camera parameters aperture (Ap), shutter speed (T_{Im}), and sensor gain (ISO).

The term $\min\{1, x\} := (\min\{1, x_i\})_{i=1,2,3}$ truncates the resulting color vector to the interval $[0, 1]^3$, modeling overexposure of the camera sensor.

Model F3

We separately model the effect of the polarizers on each of the frequencies produced by the light source. Our model becomes

$$\begin{pmatrix} \tilde{R} \\ \tilde{G} \\ \tilde{B} \end{pmatrix} = \min\{1, eWS\operatorname{diag}\left((T^p - T^c)\cos^2(\theta_1 - \theta_2) + T^c\right)\frac{1}{255}\begin{pmatrix} R \\ G \\ B \end{pmatrix}\}, \tag{12}$$

where $T^p := (T_R^p, T_G^p, T_B^p) \in [0, 1]^3$ are the transmission rates for each color when the polarizers are aligned, and $T^c \in [0, 1]^3$ when their polarization axes are perpendicular. From the technical specifications of the light source and polarizers, appropriate values are $T^p = (0.29, 0.35, 0.33)$ and $T^c = (0.02, 0.08, 0.18)$; for the corresponding datasheets, see the `light_source` and `polarizer` components in Table 9.

Appendix V

Causal Ground Truth

In this appendix, we formalize a causal interpretation of the ground-truth graphs in Fig. 3. Equipped with this interpretation, which we give in Definition 1, the graphs describe constraints on the underlying causal system. We also provide a procedure to empirically validate these constraints using data collected from the chambers.

In Sec. 2, we categorized the variables in each ground-truth graph G as *non-manipulable* variables, corresponding to sensor measurements, and *manipulable* variables, that is, actuators and sensor parameters. Here, we further split the latter into manipulable variables with no incoming edges in G, which we call *exogenous*, and manipulable variables with incoming edges, which we call *endogenous* and are set by the control computer as a function of the other variables in the system. We refer to such functions as *assignment functions*.

Let $X_V = (X_1, \dots, X_k)$ denote a vector of chamber variables and let us assume that there is a true causal model describing the system; more formally, we denote by $P(X_j^t \mid \mathrm{do}(X_{V\setminus\{j\}}^{t_0} = x))$ the distribution of variable X_j at time t, after performing an intervention on all other chamber variables $X_{V\setminus\{j\}}$ at some earlier time $t_0 < t$. We can sample from this distribution by taking a measurement from the chamber, where the intervention corresponds to the following manipulations of its variables: if X_i is a non-manipulable variable, i.e., a sensor measurement, the intervention consists of overwriting its value with x_i. If $X_k \in X_{V\setminus\{j\}}$ is an exogenous manipulable variable, then the intervention corresponds to setting it to x_k. If it is an endogenous manipulable variable, the intervention also sets its value to x_k, overriding the value set by its assignment function.

Definition 1 (Ground-truth causal model)**.** *Let G be a directed graph over X_V. We call G a* ground-truth graph *if an edge $X_i \to X_j$ in G implies that*

$$\begin{aligned}&\exists t, T > 0, x^A, x^B \textit{ with } x_i^A \neq x_i^B \textit{ and } x_k^A = x_k^B \textit{ for all } k \neq i,\\ &\textit{such that } P(X_j^{t+T} \mid \mathrm{do}(X_{V\setminus\{j\}}^t = x^A)) \neq P(X_j^{t+T} \mid \mathrm{do}(X_{V\setminus\{j\}}^t = x^B)).\end{aligned} \tag{13}$$

Although it is impossible to directly evaluate (13)—we would need to perform two different interventions simultaneously—Definition 1 allows us, under additional assumptions, to make statements about what a ground-truth graph G implies for the observations collected from the chambers. For example, if we assume that the distribution resulting from an intervention depends on the time lag (but not the absolute time t) at which the intervention is performed (the assumption is formalized in Proposition 1 below), then we can use two-sample tests for testing (13). This assumption is suitable for the controlled environment of the light tunnel and for many of the effects in the wind tunnel. For the effects on the barometric measurements $(\tilde{P}_{\mathrm{up}}, \tilde{P}_{\mathrm{dw}}, \tilde{P}_{\mathrm{amb}}, \tilde{P}_{\mathrm{int}})$, whose distribution changes with time due to shifts in atmospheric pressure, the assumption is still reasonable for measurements taken closely in time.

V.1 Empirical validation

According to Definition 1, to validate an effect $X_i \to X_j$, we would need to reject the null hypothesis

$$\begin{aligned}H_0 :&\forall t, T > 0, x^A, x^B \text{ with } x_i^A \neq x_i^B \text{ and } x_k^A = x_k^B \text{ for } k \neq i,\\ &\text{such that } P(X_j^{t+T} \mid \mathrm{do}(X_{V\setminus\{j\}}^t = x^A)) = P(X_j^{t+T} \mid \mathrm{do}(X_{V\setminus\{j\}}^t = x^B)).\end{aligned} \tag{14}$$

However, as discussed after Definition 1, additional assumptions are needed to test H_0, since we cannot perform the two different interventions $\mathrm{do}(X_{V\setminus\{j\}}^t = x^A)$ and $\mathrm{do}(X_{V\setminus\{j\}}^t = x^B)$ at the same time t. We first state our validation procedure and then provide sufficient conditions for it to have the appropriate level for testing H_0.

Validation procedure Pick predefined x_A, x_B, T, N, α.

1. sample $U^1, \dots, U^N$ i.i.d. from a Bernoulli distribution with $p = 0.5$
2. sample $\Delta t_1, \dots, \Delta t_N$ i.i.d. $\sim \mathrm{Unif}[10^{-3}, 1]$
3. let $t_0 + T$ be the current time
4. for $n = 1, \dots, N$
   (a) wait Δt_n seconds
   (b) (at time $t_n := t_{n-1} + T + \Delta t_n$) set $X_{V\setminus\{j\}} \leftarrow x^A$ if $U^n = 0$, or $X_{V\setminus\{j\}} \leftarrow x^B$ otherwise
   (c) wait T seconds

   (d) (at time $t_n + T$) collect the measurement x^{t_n+T} from the chamber

   (e) add the measurement of X_j to the sample $\boldsymbol{X}^A \leftarrow \boldsymbol{X}^A \cup \{x_j^{t_n+T}\}$ if $U^n = 0$, and $\boldsymbol{X}^B \leftarrow \boldsymbol{X}^B \cup \{x_j^{t_n+T}\}$ otherwise

5. perform a two-sample test ψ that is level α (in practice, we use a Kolmogorov-Smirnov [6] test), comparing the samples $\boldsymbol{X}^A$ and $\boldsymbol{X}^B$, resulting in a p-value p

6. reject H_0 if $p \leq \alpha$

The above procedure amounts to a randomized controlled trial, where the treatment units are the N measurements, and both the intervention allocations and intervention times are randomized. Proposition 1 provides sufficient conditions for the procedure to achieve the correct level in testing H_0 (14).

Proposition 1. *Let x_A, x_B, T, N, α be the values used in the validation procedure and assume that the following two assumptions hold.*

A1. *For all $t_1, \ldots, t_N$ such that $t_i + T < t_{i+1}$ for all $i \in \{1, \ldots, N-1\}$, and for all $x^1, \ldots, x^N \in \{x_A, x_B\}$, it holds that for all $n \in \{1, \ldots, N\}$*

$$P(X_j^{t_n+T} \mid \mathrm{do}(X_{V\setminus\{j\}}^{t_n} = x^n), \ldots, \mathrm{do}(X_{V\setminus\{j\}}^{t_1} = x^1), X_j^{t_{n-1}+T}, \ldots, X_j^{t_1+T}) = P(X_j^{t_n+T} \mid \mathrm{do}(X_{V\setminus\{j\}}^{t_n} = x^n)).$$

A2. $\forall \tau, t$

$$P(X_j^{t+T} \mid \mathrm{do}(X_{V\setminus\{j\}}^{t} = x_A)) = P(X_j^{t+T+\tau} \mid \mathrm{do}(X_{V\setminus\{j\}}^{t+\tau} = x_A)),$$

and

$$P(X_j^{t+T} \mid \mathrm{do}(X_{V\setminus\{j\}}^{t} = x_B)) = P(X_j^{t+T+\tau} \mid \mathrm{do}(X_{V\setminus\{j\}}^{t+\tau} = x_B)).$$

Then, the validation procedure is level α for testing H_0.

Proof. Assumptions A1 and A2 imply that there is P^A and P^B such that, conditioned on the realizations of $U_1, \ldots, U_N$, $\boldsymbol{X}^A \sim$ i.i.d. P^A, and $\boldsymbol{X}^B \sim$ i.i.d. P^B. Additionally, under H_0, $P^A = P^B$. Thus, $P(\psi(\boldsymbol{X}^A, \boldsymbol{X}^B) = 1) \leq \alpha$. ☐

In essence, assumptions *A1* and *A2* ensure that the test in step 5 of the procedure receives two samples composed of independent and identically distributed observations. Alternative tests requiring weaker conditions may be possible. In any case, the assumptions are not unreasonable for the causal chambers. Both assumptions hold when X_j is an endogenous manipulable variable, and they are reasonable for all sensor measurements in the light tunnel, and most in the wind tunnel. An exception comes from the barometric measurements $\tilde{P}_{\mathrm{dw}}, \tilde{P}_{\mathrm{up}}, \tilde{P}_{\mathrm{amb}}$ and $\tilde{P}_{\mathrm{int}}$, which are affected by the atmospheric pressure outside the chamber. In this case, assumption *A2* may not hold for large τ. However, the issue is largely reduced for measurements taken very closely in time. Furthermore, the randomization of time points and intervention assignments provides some robustness against violations of A1 and A2.

The results of the procedure, in the form of the p-values computed in step 5 (using the Kolmogorov-Smirnov two-sample test), are shown for each edge and ground-truth graph in Table 5-Table 8. The corresponding datasets are listed in the caption of each table, and the code to compute the p-values can be found in the Jupyter notebook `causal_validation.ipynb` in the paper repository at `github.com/juangamella/causal-chamber-paper`. The data from some experiments was used to compute the p-values for different edges and some experiments have been repeated a few times; in Table 5-Table 8 we show the raw p-values produced by the test without additional corrections for multiple testing.

Edge	p-value	Edge	p-value	Edge	p-value	Edge	p-value
$R \to \tilde{I}_1$	2.3×10^{-14}	$G \to \tilde{I}_1$	1.6×10^{-14}	$B \to \tilde{I}_1$	1.9×10^{-14}	$R \to \tilde{I}_2$	2.3×10^{-14}
$G \to \tilde{I}_2$	1.6×10^{-14}	$B \to \tilde{I}_2$	1.9×10^{-14}	$R \to \tilde{I}_3$	2.3×10^{-14}	$G \to \tilde{I}_3$	1.6×10^{-14}
$B \to \tilde{I}_3$	1.9×10^{-14}	$R \to \tilde{V}_1$	2.3×10^{-14}	$G \to \tilde{V}_1$	1.6×10^{-14}	$B \to \tilde{V}_1$	1.9×10^{-14}
$R \to \tilde{V}_2$	2.3×10^{-14}	$G \to \tilde{V}_2$	1.6×10^{-14}	$B \to \tilde{V}_2$	1.9×10^{-14}	$R \to \tilde{V}_3$	2.3×10^{-14}
$G \to \tilde{V}_3$	1.6×10^{-14}	$B \to \tilde{V}_3$	1.9×10^{-14}	$R \to \tilde{C}$	2.3×10^{-14}	$G \to \tilde{C}$	1.6×10^{-14}
$B \to \tilde{C}$	1.9×10^{-14}	$\theta_1 \to \tilde{I}_3$	3×10^{-14}	$\theta_2 \to \tilde{I}_3$	6.6×10^{-14}	$\theta_1 \to \tilde{V}_3$	3×10^{-14}
$\theta_2 \to \tilde{V}_3$	6.6×10^{-14}	$\theta_1 \to \tilde{\theta}_1$	3×10^{-14}	$\theta_2 \to \tilde{\theta}_2$	6.6×10^{-14}	$R_1 \to \tilde{\theta}_1$	1.6×10^{-14}
$O_1 \to \tilde{\theta}_1$	8.3×10^{-101}	$R_2 \to \tilde{\theta}_2$	2.3×10^{-14}	$O_2 \to \tilde{\theta}_2$	3.4×10^{-67}	$R_C \to \tilde{C}$	1.9×10^{-14}
$O_C \to \tilde{C}$	6×10^{-24}	$L_{11} \to \tilde{I}_1$	1.6×10^{-14}	$L_{12} \to \tilde{I}_1$	2.3×10^{-14}	$L_{11} \to \tilde{V}_1$	1.6×10^{-14}
$L_{12} \to \tilde{V}_1$	2.3×10^{-14}	$T_1^I \to \tilde{I}_1$	1.9×10^{-14}	$D_1^I \to \tilde{I}_1$	3×10^{-14}	$T_1^V \to \tilde{V}_1$	1.9×10^{-14}
$D_1^V \to \tilde{V}_1$	1.9×10^{-14}	$L_{21} \to \tilde{I}_2$	4.2×10^{-14}	$L_{22} \to \tilde{I}_2$	6.6×10^{-14}	$L_{21} \to \tilde{V}_2$	4.2×10^{-14}
$L_{22} \to \tilde{V}_2$	1.3×10^{-12}	$T_2^I \to \tilde{I}_2$	1.9×10^{-14}	$D_2^I \to \tilde{I}_2$	1.6×10^{-14}	$T_2^V \to \tilde{V}_2$	1.6×10^{-14}
$D_2^V \to \tilde{V}_2$	8.9×10^{-13}	$L_{31} \to \tilde{I}_3$	2.3×10^{-14}	$L_{32} \to \tilde{I}_3$	4.2×10^{-14}	$L_{31} \to \tilde{V}_3$	6.8×10^{-12}
$L_{32} \to \tilde{V}_3$	4.2×10^{-14}	$T_3^I \to \tilde{I}_3$	2.3×10^{-14}	$D_3^I \to \tilde{I}_3$	4.2×10^{-14}	$T_3^V \to \tilde{V}_3$	1.9×10^{-14}
$D_3^V \to \tilde{V}_3$	2.3×10^{-14}						

Table 5: Results of the validation procedure for the edges of the standard-configuration graph of the light tunnel (Fig. 3a). For each edge, we show the p-value resulting from the Kolmogorov-Smirnov [6] test computed in step 5 of the procedure. All p-values are below 10^{-11}. The experimental data, together with the interventions x^A, x^B, waiting time T and sample size N for each edge can be found in the `lt_validate_v1` dataset at `causalchamber.org`.

Edge	p-value	Edge	p-value	Edge	p-value
$L_{\text{in}} \to \tilde{\omega}_{\text{in}}$	1.6×10^{-14}	$T_{\text{in}} \to \tilde{\omega}_{\text{in}}$	1.4×10^{-152}	$L_{\text{in}} \to \tilde{\omega}_{\text{out}}$	1.6×10^{-14}
$L_{\text{in}} \to \tilde{C}_{\text{in}}$	1.6×10^{-14}	$L_{\text{in}} \to \tilde{C}_{\text{out}}$	1.8×10^{-22}	$L_{\text{out}} \to \tilde{\omega}_{\text{in}}$	3.4×10^{-14}
$L_{\text{out}} \to \tilde{\omega}_{\text{out}}$	3.4×10^{-14}	$T_{\text{out}} \to \tilde{\omega}_{\text{out}}$	8.3×10^{-122}	$L_{\text{out}} \to \tilde{C}_{\text{out}}$	3.4×10^{-14}
$L_{\text{out}} \to \tilde{C}_{\text{in}}$	1.1×10^{-8}	$H \to \tilde{\omega}_{\text{in}}$	1.5×10^{-17}	$H \to \tilde{\omega}_{\text{out}}$	4×10^{-58}
$L_{\text{in}} \to \tilde{P}_{\text{int}}$	2.5×10^{-7}	$H \to \tilde{P}_{\text{int}}$	8.8×10^{-11}	$L_{\text{out}} \to \tilde{P}_{\text{int}}$	0.082
$O_{\text{int}} \to \tilde{P}_{\text{int}}$	1.1×10^{-23}	$L_{\text{in}} \to \tilde{P}_{\text{up}}$	1.6×10^{-14}	$H \to \tilde{P}_{\text{up}}$	3×10^{-14}
$L_{\text{out}} \to \tilde{P}_{\text{up}}$	3.4×10^{-14}	$O_{\text{up}} \to \tilde{P}_{\text{up}}$	8×10^{-7}	$L_{\text{in}} \to \tilde{P}_{\text{dw}}$	1.6×10^{-14}
$H \to \tilde{P}_{\text{dw}}$	3×10^{-14}	$L_{\text{out}} \to \tilde{P}_{\text{dw}}$	3.4×10^{-14}	$O_{\text{dw}} \to \tilde{P}_{\text{dw}}$	1.7×10^{-5}
$O_{\text{amb}} \to \tilde{P}_{\text{amb}}$	1.5×10^{-284}	$O_{\text{in}} \to \tilde{C}_{\text{in}}$	9.2×10^{-22}	$R_{\text{in}} \to \tilde{C}_{\text{in}}$	2.8×10^{-59}
$O_{\text{out}} \to \tilde{C}_{\text{out}}$	2.5×10^{-22}	$R_{\text{out}} \to \tilde{C}_{\text{out}}$	5×10^{-59}	$A_1 \to \tilde{S}_1$	2.2×10^{-59}
$O_1 \to \tilde{S}_1$	6.9×10^{-36}	$R_1 \to \tilde{S}_1$	2.8×10^{-59}	$A_1 \to \tilde{S}_2$	1.1×10^{-55}
$A_2 \to \tilde{S}_2$	2.2×10^{-59}	$O_2 \to \tilde{S}_2$	9.5×10^{-35}	$R_2 \to \tilde{S}_2$	6×10^{-59}
$A_1 \to \tilde{M}$	4.5×10^{-9}	$L_{\text{in}} \to \tilde{M}$	1.8×10^{-9}	$L_{\text{out}} \to \tilde{M}$	2.7×10^{-7}
$H \to \tilde{M}$	5.7×10^{-16}	$O_M \to \tilde{M}$	1×10^{-32}	$R_M \to \tilde{M}$	4.2×10^{-59}

Table 6: Results of the validation procedure for the edges of the standard-configuration graph of the wind tunnel (Fig. 3c). For each edge, we show the p-value resulting from the Kolmogorov-Smirnov [6] test computed in step 5 of the procedure. The computed p-values are all below 8×10^{-7}, except for the one corresponding to the edge $L_{\text{out}} \to \tilde{P}_{\text{int}}$ (0.082). The data to validate this edge was collected during a windy day (16.4.2024 in Zurich), and the effect on the intake barometer (a few pascals) was small compared to the large fluctuations (100 pascals) in the ambient atmospheric pressure. When using a more powerful rank-sum test [7], we obtained a p-value of 0.029. The experimental data, together with the interventions x^A, x^B, waiting time T and sample size N for each edge can be found in the `wt_validate_v1` dataset at `causalchamber.org`.

Edge	p-value	Edge	p-value	Edge	p-value	Edge	p-value
$\theta_1 \to \tilde{\text{Im}}$	2×10^{-13}	$\theta_2 \to \tilde{\text{Im}}$	1.6×10^{-14}	$R \to \tilde{\text{Im}}$	2.3×10^{-14}	$G \to \tilde{\text{Im}}$	1.6×10^{-14}
$B \to \tilde{\text{Im}}$	1.9×10^{-14}	$T_{\text{Im}} \to \tilde{\text{Im}}$	3×10^{-14}	$\text{Ap} \to \tilde{\text{Im}}$	2.3×10^{-14}	$\text{ISO} \to \tilde{\text{Im}}$	4.2×10^{-14}

Table 7: Results of the validation procedure for the additional edges in the "camera" graph of the light tunnel (Fig. 3b). For each edge, we show the p-value resulting from the Kolmogorov-Smirnov [6] test computed in step 5 of the procedure. All p-values are below 10^{-12}. For the image variable $\tilde{\text{Im}}$, the test is performed on the average of pixel values. The experimental data, together with the interventions x^A, x^B, waiting time T and sample size N for each edge can be found in the `lt_camera_validate_v1` dataset at `causalchamber.org`.

Edge	p-value
$\tilde{P}_{\text{dw}} \to L_{\text{in}}$	2.3×10^{-14}
$\tilde{P}_{\text{dw}} \to L_{\text{out}}$	2.3×10^{-14}

Table 8: Results of the validation procedure for the additional edges of the "pressure-control" graph of the wind tunnel (Fig. 3d). For each edge, we show the p-value resulting from the Kolmogorov-Smirnov [6] test computed in step 5 of the procedure. The experimental data, together with the interventions x^A, x^B, waiting time T and sample size N for each edge can be found in the `wt_pc_validate_v1` dataset at `causalchamber.org`.

Appendix VI

Component Datasheets

Component	Chamber	Description and datasheets
`light_source`	`wt/lt`	The light source of the light tunnel. - LED array: `datasheets/light_source.pdf` - Individual LED (e.g. for wavelengths): `datasheets/light_source_led.pdf`
`current_sensor`	`wt/lt`	Current sensor for the light source and fans. `datasheets/current_sensor.pdf`
`motor`	`wt/lt`	The stepper motors to control the polarizers and the hatch. - Motor: Model 17HS8401 in datasheet `datasheets/motor.pdf` - A4988 motor driver: `datasheets/motor_driver.pdf`
`angle_sensor`	`lt`	The rotary potentiometer used to measure the polarizer angle. Model 3590S-6-502L in datasheet `datasheets/angle_sensor.pdf`
`light_sensor`	`lt`	The light-intensity sensor of the light tunnel. Model Si1151-AB00-GMR in datasheet `datasheets/light_sensor.pdf`
`led`	`lt`	The LEDs placed by each light sensor in the light tunnel. - LEDs: `datasheets/led.pdf` - Digital rheostats: model MCP4151 in `datasheets/potentiometer.pdf`
`camera`	`lt`	The light-tunnel camera, a Sony α6100: `datasheets/camera.pdf`
`arduino`	`wt/lt`	The control computer used in the chambers, i.e., an Arduino Mega Rev3: `datasheets/arduino_mega.pdf`
`fan`	`wt`	The wind tunnel fans: - Fan specifications: `datasheets/fan.pdf` - Speed control white paper: `datasheets/fan_pwm.pdf` - Specifications of a similar motor: `datasheets/fan_motor.pdf`
`barometer`	`wt`	The high precision barometers used in the wind tunnel: `datasheets/barometer.pdf`
`potentiometer`	`wt`	Digital potentiometer used to regulate amplitude in the amplification circuit of the wind tunnel. It is the same component as the rheostats of the light tunnel, i.e., model MCP4151 in `datasheets/potentiometer.pdf`
`microphone`	`wt`	The wind tunnel microphone consists of an electret microphone with the LM385-DIP8 amplifier in `datasheets/microphone_amplifier.pdf`
`speaker`	`wt`	The wind-tunnel speaker. - Amplifier: model LM386-D08-T in `datasheets/speaker_amplifier.pdf` - Speaker: `datasheets/speaker.pdf`

Table 9: Physical components of the chambers and their related datasheets. The datasheets can be found at `github.com/juangamella/causal-chamber/`.

References

[1] Yasuki Nakayama. *Introduction to fluid mechanics.* Butterworth-Heinemann, 2018.

[2] A. Kayode Coker. Ludwig's applied process design for chemical and petrochemical plants volume 1. *Chemical Engineering*, 115(7):8–9, 2008.

[3] Johann Tang. Fan basics: Air flow, static pressure, and impedance. `https://blog.orientalmotor.com/fan-basics-air-flow-static-pressure-impedance`, 2019.

[4] Edward Collett. *Field guide to polarization.* International society for optics and photonics, 2005.

[5] Jose Lages, Remo Giust, and Jean-Marie Vigoureux. Composition law for polarizers. *Physical Review A*, 78(3): 033810, 2008.

[6] Nickolay Smirnov. Table for estimating the goodness of fit of empirical distributions. *The Annals of Mathematical Statistics*, 19(2):279–281, 1948.

[7] Henry B. Mann and Donald R. Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics*, pages 50–60, 1947.