

A Method for Constraint Inference Using Pose and Wrench Measurements

Guru Subramani, Michael Hagenow, Michael Gleicher, *Member, IEEE* and Michael Zinn, *Member, IEEE*

Abstract—Many physical tasks such as pulling out a drawer or wiping a table can be modeled with geometric constraints. These geometric constraints are characterized by restrictions on kinematic trajectories and reaction wrenches (forces and moments) of objects under the influence of the constraint. This paper presents a method to infer geometric constraints involving unmodeled objects in human demonstrations using both kinematic and wrench measurements. Our approach takes a recording of a human demonstration and determines what constraints are present, when they occur, and their parameters (e.g. positions). By using both kinematic and wrench information, our methods are able to reliably identify a variety of constraint types, even if the constraints only exist for short durations within the demonstration. We present a systematic approach to fitting arbitrary scleronomic constraint models to kinematic and wrench measurements. Reaction forces are estimated from measurements by removing friction. Position, orientation, force, and moment error metrics are developed to provide systematic comparison between constraint models. By conducting a user study, we show that our methods can reliably identify constraints in realistic situations and confirm the value of including forces and moments in the model regression and selection process.

Index Terms—Geometric constraints, kinematics, teaching by demonstration

I. INTRODUCTION

TASKS performed by robots frequently include constraints. For example, a manipulator is restricted to linear motion when pulling out a drawer. Representing physical interactions as constraints provides information about the permissible motion of objects in the environment and the permissible directions that forces and moments may be applied to them. This knowledge provides a way to inform the robot to restrict its movements to allowable configurations so it does not violate this closed kinematic chain (e.g., to plan motion consistent with the constraint [1]–[3]) and to apply safe levels of force against the constraint to achieve the appropriate task-specific interaction (e.g., through hybrid control [4]).

Information about the constraints in the task is valuable to program a robot, but current approaches to determine constraint parameters are not straightforward. For example, commercial robotic systems provide hybrid control abilities, however the constraints and their geometry must be determined manually. This makes such control difficult to use, limiting the ability of such systems to perform tasks involving constraints. One solution is to infer constraints from a human demonstration, which eliminates the need for explicit programming [5], [6]. This paper presents an approach to




Fig. 1: Top: Demonstration of an espresso making task using instrumented tongs consisting of sliding an espresso cup, pulling out a drawer, and actuating an espresso lever. Center: Inferred constraints in task. Our approach infers the samples associated with each constraint (segmentation), the type of constraint (model selection) and model parameters of each constraint (model regression) in the task.

infer constraints in human demonstrations of recorded force, moment, position, and orientation measurements.

The requirements of a constraint inference system may vary dramatically depending on application. In manufacturing environments, the tool and environment geometry may be known beforehand, but other applications may require interacting with unknown objects. If a system is to perform inference using a natural demonstration (e.g the espresso making task in Fig. 1), it may also need to handle demonstrations involving multiple constraint interactions, short constraint interactions, and other physical interactions such as picking up and placing objects.

Past work, summarized below, has shown that geometric constraint models can be fit from a demonstration containing sufficient data, however limited attention has been given to the ability to recognize constraints in a short demonstration. For example, the method described in [7] effectively learns

a constraint model and develops a policy to act in the null-space of a constraint, which allows mapping of motion across constraints. To accomplish this, the constraint model is trained using 20 seconds of recorded demonstration data. Particularly when a task is performed in a natural way, there is no guarantee a model can be fit with the available duration of demonstration data. Certain applications might require a constraint inference method that minimizes the data required for fitting.

Representations of constraints can be general (i.e., a broad model that fits many types of constrained motion) [1], [8] or can be derived from geometric primitives [9]. A general model offers flexibility to learn a wider range of constraints without specification of additional models, but is difficult to interpret. On the other hand, primitive models are more easily understood by human compatriots. For example, a rotational motion can be interpreted by an operator, whereas a general model or combination of low-level restrictions on rotational or translational motion may have little semantic meaning. Great effort [10]–[13] has been devoted to constraint fitting with known object geometries (e.g., polyhedral, cylindrical, and spherical tools) by combining kinematic and wrench relationships to fit the constraints of the interaction, including simultaneous active constraint interactions. Similarly, our method applies wrench-based fitting to our previous methods [14], [15] using nonlinear regression of constraint parameters, which allow for parameterization even when the interacting objects are not known.

The contribution of this paper is in presenting and assessing a constraint inference approach that uses both kinematic and wrench information to infer geometric constraints in a human demonstration, even if the objects are unknown. Our approach can take a recording of a demonstration and determine what constraints are present, when they occur, and their parameters. It can reliably identify and localize a variety of constraint types, even when the constrained interactions are short. We build upon earlier work [14], [15] that uses only kinematics to fit constraint models involving unknown objects and uses forces and moment information for validation. More specifically, the new method extends our previous work by providing a systematic review of the value of wrench information throughout the constraint inference process, extending to new constraint models, and providing a more thorough and formalized explanation of the constraint inference approach. We provide new technical method details which are validated through new, extensive experimental results.

In our new method, we incorporate reaction forces and moments in an optimization objective to fit constraint models. We describe an approach to resolve reaction wrenches from wrench measurements which also contain constraint specific friction wrenches. Our method requires minimal assumptions about the friction model and does not require any specific friction parameters. We define a general set of position, force, orientation, and moment error metrics that are used to compare fit models with measured data. The product of these contributions is a constraint inference approach that explicitly recognizes the physical interaction that occurs when interacting with a constraint, where motions and wrenches are




Fig. 2: A User demonstrates a chess move by sliding the knight against the chess board. Fitting the *planar-relaxed constraint* using kinematic information [red] is compared against fitting with a combination of pose and wrench measurements (Kinematic-Wrench fitting) [green]. Kinematic-Wrench fitting is able to identify the plane using considerably fewer samples than kinematic fitting.

related in a manner dictated by the physics of the constrained interaction. Our approach takes advantage of both wrench and kinematic measurements which enables improved performance over our previous methods that use only kinematics to fit constrained interactions with unknown objects. We assess our approach with a human subjects study with 9 participants that confirms both the approach’s performance as well as the value of including wrench information directly into the nonlinear fitting of constraints, even if there are unknown objects.

II. CONSTRAINT INFERENCE OVERVIEW

This paper provides a framework to infer geometric constraints in human demonstrations. Our approach is to first partition the human demonstration into constrained motion and unconstrained motion, fit all candidate constraint models to each constrained motion segment, and finally select the appropriate constraint models from the fit models. We organize the constraint inference as three sequential steps:

- 1) **Demonstration segmentation:** Applications may require determining when constraints occur in a demonstration, especially when demonstrations represent a larger task. We discuss methods to segment the constrained motion from demonstrations in Section VI (after model regression and model selection to cover the more important contributions first). Our approach takes advantage of the reaction forces and moments that exist during constraint interaction.
- 2) **Model regression:** In Section IV, we show how nonlinear least squares regression can be used to fit geometric constraint models to kinematic and wrench measurements by leveraging the principle of virtual work. Our method derives reaction forces and moments from mea-

sured wrenches by considering contributing dynamics from friction. This is described in Section III-C. We show that by simultaneously fitting with kinematic and wrench data, we are able to fit certain constraints using significantly less data compared to fitting with kinematic data (see Fig. 2).

- 3) **Model selection:** Once all models are fit to the constrained segments, each model fit must be evaluated and the appropriate model must be selected. The regression objectives are inadequate to perform this comparison because they evaluate to different units and scales. To address this problem, kinematic and wrench error metrics are derived for the general scleronomic constraint model to provide a consistent approach to evaluate models. This is described in Section V.

III. MATHEMATICAL MODELING OF CONSTRAINTS

Our constraint inference approach requires defined mathematical representations of the constraints. We use models based on human interpretable geometric primitives to represent constrained motion such as the *planar constraint* (e.g. a flat object wiping against a table), *point-on-plane constraint* (e.g. the motion of a pen on paper), and *axial rotation constraint* (e.g. the motion of a hinged door). These constraints are represented as mathematical equations that the tool or object must satisfy when under the influence of the constraint. The constraint models and relationship between the reaction wrenches (reaction forces and moments) imparted on constrained objects is described below.

When a robot end effector or other object moves under the influence of a constraint, the degrees of freedom of its rigid body motion are restricted to a subset in $SE(3)$. These restrictions may be represented as mathematical relationships between defined geometry on the body (e.g., a point on the body) and defined global geometry (e.g., a plane in the environment). For example, consider a pen moving against paper. The motion of the pen tip is constrained to the planar surface of the paper. In this example, the rigid body is the pen, the location of the pen tip is the defined geometry on the rigid body, and the planar surface represents the global geometry. Geometric constraint models consist of equations that define the possible configurations of the rigid body.

Geometric constraints impose restrictions on the configurations on the rigid body. Applying forces and moments on a constrained rigid body creates reaction forces and moments between the rigid body and the constraint. The principle of virtual work expresses that reaction forces and moments from the constraint do not produce work. This property can be used to determine the permissible reaction forces and reaction moments of the constraint. It is also necessary to consider additional forces that may act on the rigid body such as friction, inertia, and gravity. These concepts are developed below.

A. Geometric constraint models

In our formulation, we consider scleronomic constraints, which are geometric constraint models that only impose restrictions on position and orientation (i.e. generalized coordinates of the rigid body) and do not vary with time [16]. In this paper we consider the eight constraint models

shown in Fig. 3. These constraints were chosen as examples of commonly occurring geometric primitives. Our approach can be extended to other models satisfying the scleronomic constraint formulation. The constraint models considered here include:

- **Point-on-line constraint:** The *point-on-line constraint* is equivalent to drawing a line against a straight edge. A point on the rigid body is constrained to a line in space.
- **Planar constraint:** The *planar constraint* is exemplified by an eraser moving against a whiteboard. The rigid body can only rotate about a vector perpendicular to the plane, and all points within the rigid body translate parallel to this plane.
- **Prismatic constraint:** The *prismatic constraint* represents translational motion in one direction. It is similar to pulling out a drawer. All points on the rigid body translate identically.
- **Axial rotation constraint** The *axial rotation constraint* is similar to a door knob or a hinged door. All points on the rigid body rotate about an axis and translations are not permitted. This constraint is also referred to as a revolute joint.
- **Relaxed constraints:** The *planar*, *axial*, and *prismatic* constraints impose restrictions on the orientation of the rigid body. In certain cases, the demonstration tool may not rigidly fix the orientation of the manipulated rigid body. For example, when grabbing an object with a parallel gripper, there is a slip direction where the rigid body may rotate. The *relaxed* versions of these constraints remove restrictions on the orientation of the constrained object in order to model this type of behavior.

The detailed constraint mathematical models are developed in Appendix A. However, to provide context we present an extended description of the *point-on-plane constraint* development in the following section.

1) Extended Description for Point-on-plane Constraint

Consider a 6-degree of freedom rigid body of negligible inertial properties located in space through translational coordinates $\mathbf{r} \in \mathbb{R}^3$ and orientation coordinates represented by a unit quaternion $\mathbf{q} \in \text{Spin}(3)$. The measurement reference frame is the fixed frame on this rigid body. The rigid body is constrained by k constraint equations $\Phi : SE(3) \Rightarrow \mathbb{R}^k$:

$$\Phi(\mathbf{r}, \mathbf{q}) = 0 \quad (1)$$

The scleronomic constraint model represented by equation (1) may be instantiated with a specific constraint model (e.g. with a *point-on-plane constraint*) with associated model parameters $\boldsymbol{\alpha}$ which configure the constraint (e.g. the position and orientation of the plane). Thus, Φ becomes $\Phi(\mathbf{r}, \mathbf{q}, \boldsymbol{\alpha})$.

The *point-on-plane constraint* describes a point on the rigid body constrained to a plane in the environment such as a pencil tip (point) moving across paper (plane). An illustrated representation of this constraint is shown in Fig. 4.

A plane may be specified by applying a general displacement (i.e. translation and rotation) transformation of the x - y plane which involves:




Fig. 3: Geometric constraints considered in this paper. The three constraint models *planar constraint*, *prismatic constraint* and *axial constraint* enforce constraints on the orientation of the constrained object. The *relaxed* counterparts are similar constraint models except with removed orientation constraints. The *point-on-plane* and *point-on-line* constraint models contain representations of the location of contact with respect to the measured position on the object. The measurement frame is shown in red.




Fig. 4: *Point-on-plane constraint* model. The point s^* is constrained to the planar surface. $\bar{s}$ is a fixed vector from the origin of the measurement frame to point s^* (expressed in the measurement frame). The global x-y plane is transformed using the exponential map, $e^{\tilde{w}}$, to create the planar surface.

- 1) translating the x-y coordinate plane along the z-axis
- 2) rotating the translated plane about the origin.

We represent the rotation transformation using two exponential coordinates $\mathbf{w} = [w_x, w_y, 0] \in \mathbb{R}^2$ corresponding to $e^{\tilde{w}} \in SO(3)$, the exponential map, which is equivalent to a rotation matrix with an axis of rotation in the x-y plane. Rodrigues' rotation formula [17] is used to compute $e^{\tilde{w}}$. $\tilde{w}$ is the skew symmetric matrix similar to equation (62) in Appendix B. The third term of $\mathbf{w}$ is zero because rotations about the z - axis (perpendicular to the plane) do not alter the plane's geometry. The translation is represented by $d \in \mathbb{R}$. The normal vector on this plane is represented by $e^{\tilde{w}}[0 \ 0 \ 1]^T$ and the shifted origin of the x-y plane is represented by $e^{\tilde{w}}[0 \ 0 \ d]^T$.

A point $\mathbf{P}$ on the plane satisfies the following equation:

$$\left(e^{\tilde{w}}[0 \ 0 \ d]^T - \mathbf{P} \right)^T e^{\tilde{w}}[0 \ 0 \ 1]^T = 0 \quad (2)$$

which specifies that the dot product between a vector within the plane and the plane normal is zero.

Consider a point $s^* \in \mathbb{R}^3$ (in the global reference frame) defined as a point fixed to the rigid body and constrained to the plane. $\mathbf{s}$ is a vector directed from the origin of the measurement reference frame to the point s^* (vector defined in the global reference frame). The local reference frame counterpart of $\mathbf{s}$ is $\bar{s}$ (whose values do not change with the motion of the measurement frame):

$$\mathbf{s}^* \equiv \mathbf{r} + \mathbf{s} \equiv \mathbf{r} + \mathbf{A}(\mathbf{q})\bar{s} \quad (3)$$

where $\mathbf{A}(\mathbf{q})$ is the orthogonal rotation matrix equivalent to the rotation represented by $\mathbf{q}$.

The constraint is satisfied when $\mathbf{s}^*$ is constrained to the plane, i.e. equation $\Phi : SE(3) \Rightarrow \mathbb{R}^1$ is:

$$\Phi_1 \equiv \left(e^{\tilde{w}}[0 \ 0 \ d]^T - \mathbf{r} - \mathbf{A}(\mathbf{q})\bar{s} \right)^T e^{\tilde{w}}[0 \ 0 \ 1]^T = 0 \quad (4)$$

The parameters of this constraint are $\alpha = (\bar{s}, d, w_x, w_y)$. The derivation of the mathematical models of the other constraints shown in Fig. 3 is given in Appendix A.

B. Reaction forces and moments

In addition to the constraint models, our method requires that we develop a relationship between reaction wrenches and the constraint models so that all measured data can be used in the constraint fitting process. The formulation used is standard in multi-body dynamics literature [18] and is provided here as a review.

We determine permissible reaction forces and moments using virtual displacements and the principle of virtual work.

The first partial derivatives of constraint equations (1) are given by:

$$\delta\Phi = \Phi_r \delta\mathbf{r} + \Phi_q \delta\mathbf{q} = 0 \quad (5)$$

where Φ_r and Φ_q are the partial derivatives of (1) and $(\delta\mathbf{r}, \delta\mathbf{q})$ are the virtual displacements. Equation (5) may also be written using virtual rotation variable $\delta\boldsymbol{\pi}$ [18]. The virtual rotation variable is related to the angular velocity of a rigid body. This relationship is similar to how the virtual displacement $\delta\mathbf{p}$ is related to linear velocity.

$$\delta\Phi = \Phi_r \delta\mathbf{r} + \Phi_\pi \delta\boldsymbol{\pi} = 0 \quad (6)$$

Equation (6) is similar to the Pfaffian constraint form used in the robotics literature [19]. The permissible reaction forces $\mathbf{F}_r \in \mathbb{R}^3$ and reaction moments $\mathbf{N}_r \in \mathbb{R}^3$ that the constraint applies to the constrained body must satisfy the virtual work equation [20] because constraint reaction forces and moments do not produce work. The principle of virtual work requires:

$$\delta\mathbf{r}^T \mathbf{F}_r + \delta\boldsymbol{\pi}^T \mathbf{N}_r = 0 \quad (7)$$

Combining equations (6) and (7) using the Lagrange multiplier theorem [18]:

$$\Phi_r^T \boldsymbol{\lambda} + \mathbf{F}_r = 0 \quad (8)$$

$$\Phi_\pi^T \boldsymbol{\lambda} + \mathbf{N}_r = 0 \quad (9)$$

where $\boldsymbol{\lambda} \in \mathbb{R}^k$ are the Lagrange multipliers for the k constraint equations. Equations (8) and (9) provide a relationship between the constraint equations Φ and the permissible reaction forces and moments. $\Phi_\pi \delta\boldsymbol{\pi}$ is computed from $\Phi_q \delta\mathbf{q}$. This is provided in Appendix B.

C. Resolving reaction forces and moments from wrench measurements

Equations (8) and (9) relate the reaction forces and moments with the constraint equations described by equation (1). One challenge is that reaction forces and moments are not measured directly. Interactions with the physical environment involve additional forces such as friction, inertia, gravity, forces from object compliance, and forces from task-specific impedances. To fit the constraint models, we need to estimate the reaction wrenches from measurement wrenches containing these additional effects.

To do so, we assume quasi-static motion and choose to model only friction in our analysis of constrained motion. In cases where demonstrations are highly dynamic, such as involving objects with large inertia and acceleration, this approach may not be directly applicable and such dynamic effects must be modeled and incorporated appropriately.

Quasi-static motion against constraints generates frictional forces due to the sliding contact between surfaces. Forces attributed to friction oppose the motion of the constrained object. For example, when writing on paper (*point-on-plane constraint*), friction force opposes the motion of the pen tip. Friction forces may manifest themselves differently for each constraint and there is no general way to describe them.

However, friction forces and moments can be modeled for each constraint with certain assumptions. Our basic intuition is to remove friction in the direction of motion of objects in contact. We derive a relationship between the reaction

forces and moments with the measured forces and moments as functions of the parameters of the constraint:

$$\mathbf{F}_r = \mathbf{F}_r(\mathbf{r}, \mathbf{q}, \mathbf{v}, \boldsymbol{\omega}, \mathbf{F}, \mathbf{N}, \boldsymbol{\alpha}) \quad (10)$$

$$\mathbf{N}_r = \mathbf{N}_r(\mathbf{r}, \mathbf{q}, \mathbf{v}, \boldsymbol{\omega}, \mathbf{F}, \mathbf{N}, \boldsymbol{\alpha}) \quad (11)$$

where the measured quantities are forces ($\mathbf{F}$), moments ($\mathbf{N}$), linear velocity ($\mathbf{v}$), and angular velocity ($\boldsymbol{\omega}$) and the parameters of the constraint are $\boldsymbol{\alpha}$. The specific formulation of equations (10) and (11) for each constraint model is provided in Appendix A. The resolving of reaction forces and moments for the *point-on-plane* constraint is provided here as an example.

1) Resolving reaction forces and moments for the point-on-plane constraint

The goal is to resolve the reaction wrenches from the measured wrenches (i.e. determine expressions for $\mathbf{F}_r(\mathbf{r}, \mathbf{q}, \mathbf{v}, \boldsymbol{\omega}, \mathbf{F}, \mathbf{N}, \boldsymbol{\alpha})$ and $\mathbf{N}_r(\mathbf{r}, \mathbf{q}, \mathbf{v}, \boldsymbol{\omega}, \mathbf{F}, \mathbf{N}, \boldsymbol{\alpha})$). We assume that inertia of the constrained object is negligible and the only wrenches acting on the object are friction and reaction wrenches.

The force and moment balance equation for the rigid body at the measurement frame are as follows:

$$\mathbf{F} = \mathbf{F}_\mu + \mathbf{F}_r \quad (12)$$

$$\mathbf{N} = \mathbf{N}_\mu + \mathbf{N}_r \quad (13)$$

where $\mathbf{F}$ is the measured force, $\mathbf{F}_\mu$ is the friction force, $\mathbf{F}_r$ is the reaction force, $\mathbf{N}$ is the measured moment, $\mathbf{N}_\mu$ is the friction moment, and $\mathbf{N}_r$ is the reaction moment acting at the measurement frame. The measured force and moments are expressed in the global frame (i.e., rotated from the original force torque sensor measurements).

For the *point on plane constraint* seen in Fig. 4, point $\mathbf{s}^*$ on the rigid body is in contact with the plane. $\mathbf{s}^*$ moves with velocity $\mathbf{v}_s$. The friction force acts at point $\mathbf{s}^*$ on the rigid body and is directed along $\mathbf{v}_s$ ($\hat{\mathbf{v}}_s$ is a unit vector in this direction). Any force in the direction of $\mathbf{v}_s$ is assumed to be friction. $\mathbf{F}_\mu$ is the friction force located at point $\mathbf{s}^*$ directed along $\mathbf{v}_s$:

$$\mathbf{F}_\mu = (\mathbf{F} \cdot \hat{\mathbf{v}}_s) \hat{\mathbf{v}}_s \quad (14)$$

where $\mathbf{v}_s$ is given by:

$$\mathbf{v}_s = \mathbf{v} + \boldsymbol{\omega} \times \mathbf{s} \quad (15)$$

and $\mathbf{s}$ is the vector from the measurement frame to point $\mathbf{s}^*$ in the global frame.

We assume that at $\mathbf{s}^*$, no moment is transferred (including friction moment) between the rigid body and the plane. The moment induced by the friction forces at $\mathbf{s}^*$ is given by:

$$\mathbf{N}_\mu = \mathbf{s} \times \mathbf{F}_\mu \quad (16)$$

Combining equations (12) through (16) results in expressions for the reaction forces and moments as a function of measured values:

$$\mathbf{F}_r = \mathbf{F} - (\mathbf{F} \cdot \hat{\mathbf{v}}_s) \hat{\mathbf{v}}_s \quad (17)$$

$$\mathbf{N}_r = \mathbf{N} - \mathbf{s} \times (\mathbf{F} \cdot \hat{\mathbf{v}}_s) \hat{\mathbf{v}}_s \quad (18)$$

With equations to resolve reaction forces and moments, Equations 8 and 9 can be evaluated using measured data.

Similar mathematical relationships between measured forces and moments and reaction forces and moments for the other constraint models are given in Appendix A.

IV. LEAST SQUARES REGRESSION OF CONSTRAINT MODELS

Using the constraint and reaction wrench equations developed in the previous sections, the constraint model parameters can be fit to the measured data. The generalized constraint equations given in (1) are used to formulate a kinematic error objective while equations (8) and (9) are used to formulate a wrench error objective. The two error objectives are cooperative and are included in a summed objective to fit recorded demonstration data to a constraint model. Let the parameters of the constraint be α , and let the i^{th} sample of data be represented with a tuple $\mathbf{p}_i = (\mathbf{r}_i, \mathbf{q}_i, \mathbf{F}_i, \mathbf{N}_i, \mathbf{v}_i, \boldsymbol{\omega}_i)$ which consists of measured poses, wrenches, and velocities estimated from the kinematic measurements. The error objectives are defined across the N samples of data $\mathbf{p} := \{\mathbf{p}_1, \dots, \mathbf{p}_N\}$. The kinematic and wrench error objectives are described below.

A. Kinematic error objective

The kinematic error objective is developed to incorporate kinematic measurements during fitting and is given by:

$$K_{\text{error}}(\alpha, \mathbf{p}, \mathbf{w}) = \sum_{i=1}^N w_i \|\Phi(\mathbf{p}_i, \alpha)\| \quad (19)$$

where $\|\square\|$ is the L(2) norm and $\mathbf{w} := \{w_1, \dots, w_N\}$ are the weights on each sample. Sample weights allow increasing the influence of certain samples over others. Some samples included in the regression may not be part of the constraint or may be caused by measurement errors and should have less influence when fitting. For the kinematic error objective, only $(\mathbf{r}, \mathbf{q})$ is required from tuple $\mathbf{p}$.

B. Wrench error objective

A wrench error objective is developed to incorporate wrench measurements during fitting. Equations (8) and (9) provide a relationship between the constraint equations and the applied forces and moments in a demonstration. To apply equations (8) and (9) for regression, the Lagrange multipliers must be determined jointly with the parameters of the constraint. This may be formulated as a weighted least squares objective:

$$W_{\text{error}}(\alpha, \mathbf{p}, \mathbf{w}) = \sum_{i=1}^N w_i \min_{\lambda} \left(\|\Phi_{\mathbf{r}}^T \lambda + \mathbf{F}_{\mathbf{r}i}\| + \|\Phi_{\boldsymbol{\pi}}^T \lambda + \mathbf{N}_{\mathbf{r}i}\| \right) \quad (20)$$

Equations (8) and (9) are linear in λ and may be solved using linear least squares regression. The least squares error of $\mathbf{A}\lambda = \mathbf{b}$ is denoted as $LSQ(\mathbf{A}, \mathbf{b})$ where $\mathbf{A}$ and $\mathbf{b}$ are given by:

$$\mathbf{A} = \begin{bmatrix} \Phi_{\mathbf{r}}^T \\ \Phi_{\boldsymbol{\pi}}^T \end{bmatrix} \quad (21)$$

and

$$\mathbf{b} = \begin{bmatrix} -\mathbf{F}_{\mathbf{r}} \\ -\mathbf{N}_{\mathbf{r}} \end{bmatrix} \quad (22)$$

where the reaction forces and moments, $\mathbf{F}_{\mathbf{r}}$ and $\mathbf{N}_{\mathbf{r}}$, are evaluated from measured forces and moments using equations (10) and (11).

Equation (20) may be reformulated as:

$$W_{\text{error}}(\alpha, \mathbf{p}, \mathbf{w}) = \sum_{i=1}^N w_i LSQ \left(\begin{bmatrix} \Phi_{\mathbf{r}}^T(\mathbf{p}_i, \alpha) \\ \Phi_{\boldsymbol{\pi}}^T(\mathbf{p}_i, \alpha) \end{bmatrix}, \begin{bmatrix} -\mathbf{F}_{\mathbf{r}i} \\ -\mathbf{N}_{\mathbf{r}i} \end{bmatrix} \right) \quad (23)$$

The reaction forces and reaction moments are jointly estimated by including their expressions in the optimization objective.

C. Maximum likelihood estimate

Using the kinematic and wrench error objective functions we can determine the most likely model parameters α given model Φ and demonstrator data $\mathbf{p}$. The maximum likelihood estimate of model parameters α based on observations $\mathbf{p}$ is $\hat{\alpha}_{MLE}$ and can be estimated as:

$$\hat{\alpha}_{MLE} = \arg \min_{\alpha} (W_{\text{error}}(\alpha, \mathbf{p}, \mathbf{w}) + K_{\text{error}}(\alpha, \mathbf{p}, \mathbf{w})) \quad (24)$$

Since the two objectives W_{error} and K_{error} are physically related through virtual work, they behave cooperatively. Experimental evaluation indicates that weighting the two objectives has little effect on the overall fit. Equation (24) can be used to estimate the constraint parameter values.

D. Robust regression

Segmentation error and non-gaussian noise from sensor measurements may cause the nonlinear least squares optimization in equation (24) to fit poorly. All of the samples may not be part of the constraint. To address these issues, we use iterative reweighted least squares (IRLS) [21] to estimate the sample weights associated with the constraint jointly with the model parameters which helps to reject outliers.

The approach initializes equal weight values to each sample. It performs an optimization using a standard non-linear optimization algorithm (in this work we use BFGS [22]) using these weights. It calculates the errors associated with each sample and down weights samples that have large error. A re-weighting step is used to normalize the weights using a regularization function to prevent aggressive down-weighting of samples so that all samples continue to have an influence on the objective. Using these new weighted values, the algorithm performs another iteration of optimization.

The process of finding the most likely value of α is given in Algorithm 1. An error convergence exit criteria is not necessarily applicable since all models are fit to each constraint segment. In this work, we choose a basic exit criteria of maximum number of iterations N . A more advanced criteria could be considered in future implementations.

To illustrate the benefit of the iterative reweighted least squares we consider the example of moving a chess piece across a chessboard (shown in Fig. 2). In this case, we intentionally segment the demonstration incorrectly so a portion of the demonstration is in free space. When fitting the model without iterative reweighted least squares (IRLS) regression, the fit plane is offset due to outliers caused by incorrect segmentation. After 100 iterations of IRLS the fit converges to the

Algorithm 1 Iterative reweighted least squares regression

```

1: procedure ROBUSTREGRESSION( $\mathbf{X}$ ,  $\mathbf{w}$ ,  $N$ )
2:   initialize( $\mathbf{w}$ ) ▷ set all weights to 1
3:   for  $j = 1 \dots N$  do ▷ number of iterations
4:      $\mathbf{w} \leftarrow \mathbf{w} / \text{sum}(\mathbf{w})$  ▷ Normalization
5:      $\boldsymbol{\alpha} \leftarrow \text{FitFunction}(\mathbf{X}, \mathbf{w})$  ▷ BFGS minimization
6:      $\mathbf{E} \leftarrow \text{ErrorFunction}(\mathbf{X}, \boldsymbol{\alpha}, \mathbf{w})$  ▷ vector of fit errors
7:      $\mathbf{w} \leftarrow 1 / (1 + \mathbf{E}_i^{0.1})$  ▷ Regularization
8:   return  $\boldsymbol{\alpha}$  ▷

```

true constraint model (see Fig. 5). With IRLS, we can estimate the constraint parameters from imperfect measurement data.




Fig. 5: Bottom: the perspective view of the demonstration in Fig. 2 of sliding the chess piece across the board (planar constraint) is shown. The demonstration is intentionally segmented incorrectly [yellow outliers]. Iterative reweighted least squares is used to robustly fit the plane in the demonstration. 1 iteration [red] is equivalent to not using a robust fitting algorithm. This leads to a plane that has incorrect orientation. 100 iterations of IRLS [green] accurately fits the model to the constraint.

V. GEOMETRIC CONSTRAINT MODEL SELECTION

Once the parameters have been fit for a set of candidate models, it is necessary to identify the model most consistent with the measured data. The difficulty of the model selection problem is a function of the number of candidate models and the uniqueness of these models.

To simplify the model comparison process, we create a single selection metric that can be used across models. In our formulation, the optimization fit error is not appropriate as the models consist of regressor functions that are not guaranteed to have consistent units. To allow for comparison, kinematic and wrench error metrics are defined in terms of measured quantities (m, rad, N, Nm). These metrics are compared to predefined thresholds and are combined into a single error metric by counting measurements that exceed these thresholds. We refer to this process as error tallying.

A. Error metrics

Position and force information allow for the construction of separate kinematic and wrench error metrics. An error metric is defined as a measure of distance between a measurement (e.g. $\mathbf{r}, \mathbf{q}, \mathbf{F}, \mathbf{N}$) and the constraint model.

1) Kinematic error metrics

Given that a constraint can enforce both a position and an orientation, there is a need for two kinematic metrics. A position distance error metric, $D_r(\Phi(\mathbf{r}^*, \mathbf{q}^*), \mathbf{r}, \mathbf{q})$, and an

orientation distance error metric, $D_q(\Phi(\mathbf{r}^*, \mathbf{q}^*), \mathbf{r}, \mathbf{q})$, are defined for a pose measurement $(\mathbf{r}, \mathbf{q}) \in SE(3)$ where $(\mathbf{r}^*, \mathbf{q}^*)$ is coincident with the constraint (satisfies Φ). Evaluation of $(\mathbf{r}^*, \mathbf{q}^*)$ consists of finding the point on the constraint that is closest in overall distance to the measured point. Determining these values is posed as a least squares regression where $(\mathbf{r}^*, \mathbf{q}^*)$ are solved simultaneously:

$$\hat{\mathbf{r}}^*, \hat{\mathbf{q}}^* = \arg \min_{\mathbf{r}^*, \mathbf{q}^*} \begin{pmatrix} \|\mathbf{r}^* - \mathbf{r}\| + \angle(\text{conj}(\mathbf{q}) \star \mathbf{q}^*) \\ + W\Phi(\mathbf{r}^*, \mathbf{q}^*)^T \Phi(\mathbf{r}^*, \mathbf{q}^*) \\ + W\|\mathbf{q}^{*T} \mathbf{q}^* - 1\|_2 \end{pmatrix} \quad (25)$$

where W is a large weight (1000) on the constraint equations to ensure that $\mathbf{r}^*$ and $\mathbf{q}^*$ are fixed to the constraint. $\angle$ is the rotation angle of a quaternion, $\star$ represents quaternion multiplication, and conj represents the conjugate of a quaternion. $\angle(\text{conj}(\mathbf{q}) \star \mathbf{q}^*)$ computes the angle between $\mathbf{q}$ and $\mathbf{q}^*$ (bounded to 180 degrees). While the position and orientation have different units and bounds, we did not add any scaling or normalization factors as for many of the constraints, the position and orientation optimize independent aspects of the constraint.

The kinematic error metrics, D_r and D_q can subsequently be computed as:

$$D_r = \|\hat{\mathbf{r}}^* - \mathbf{r}\| \quad (26)$$

$$D_q = \angle(\text{conj}(\mathbf{q}) \star \hat{\mathbf{q}}^*) \quad (27)$$

2) Wrench error metrics

Measured forces and moments provide metrics that quantify inconsistency with the constraint model, Φ . By computing the Lagrange multipliers individually for each measurement (see Section IV-B), a force and moment error can be written as the residual of the virtual work equations previously defined in equations (8) and (9). The values are normalized with respect to the measured forces and moments to remove a dependence on magnitude, resulting in the wrench error metrics, E_f and E_n , associated with the measured forces and measured moments respectively :

$$E_F = \|\Phi_F^T \lambda + \mathbf{F}\| / \|\mathbf{F}\| \quad (28)$$

$$E_N = \|\Phi_N^T \lambda + \mathbf{N}\| / \|\mathbf{N}\| \quad (29)$$

3) Combined error metric

For model selection, we desire a single metric capable of disambiguation between candidate models. A weighted sum of the individual error metrics would contain inherent bias dependent on the range of motion and level of applied forces and moments. To circumvent this, we introduce a combined error metric (CE) which is a function of the number of samples of the constrained motion that exceed predefined error thresholds given as r_{SE} , q_{SE} , f_{SE} , and n_{SE} . The combined metric considers all measured variables and is normalized such that models can be compared directly. The error thresholds, which have units corresponding to the units of the corresponding error metric, are determined by evaluating the error metrics using the numerical values of the measurement error, which is a function of the experimental setup. The combined error

metric (summed over the measured samples, n) is evaluated as:

$$CE = \sum_{n=1}^N \left(\begin{array}{l} (D_r(n) > r_{SE}) \\ + (D_q(n) > q_{SE}) \\ + (E_f(n) > f_{SE}) \\ + (E_n(n) > n_{SE}) \end{array} \right) + 1 \quad (30)$$

Section VII-C describes how the thresholds are set for each of the error metrics based on an estimate of the experimental measurement accuracy. Our combined error metric provides one method that allows combination of the different error metrics, but does require thresholds based on an estimate of system error. Future work could include a systematic comparison between this method and other methods for combining the various forms of information, including a median error-based metric or an energy-based formulation.

For each model in the set of models, i.e., $m \in M$, a likelihood metric, $\hat{L}_m$, is evaluated from the combined error metric:

$$\hat{L}_m = \frac{CE(m)^{-1}}{\sum_m CE(m)^{-1}} \quad (31)$$

Note, the initial value of the combined error metric, CE , as defined in equation (30) is set equal to one to avoid division by zero in the calculation of the likelihood metric.

VI. SEGMENTATION

A demonstration may contain a single constraint or multiple constraints of various types. Additionally, an individual constraint may be visited multiple times during a single demonstration. Finally, free space motion may occur between constraint interactions or different constraint interactions may occur consecutively, with no intermediate free-space motion.

In order to apply regression and model selection to each constraint interaction individually, the associated samples of the constraint must be determined from the demonstration. While user-specified segmentation is possible, it is desirable to make the process automatic to avoid the need for user intervention to define constraint boundaries for segmentation.

Prior work has explored the segmentation problem. One approach [15] is to fit each constraint type over a small sliding window. Similar constraint fits are then clustered over the entire demonstration. While effective, this approach is computationally expensive and limited to simple constraint models. It also requires significant tuning of heuristics by the user to achieve reliable segmentation. In another method, [23], change point detection is used to determine constraint locations and can detect transitions between constraints without free-space motion, but this method is limited to simple models such as lines and arcs.

In our method, we use an extension to the basic idea of force thresholding. If a measured force exceeds a specified threshold (e.g. twice the measurement noise of the system), then a constraint may be present. The result is a Boolean list of contiguous samples indicating possible constrained motion segments and free space motion segments. One limitation of this simple segmentation method is that free space motion must occur between each constraint interaction. While




Fig. 6: Instrumented tongs - a handheld device used to record human demonstrations.

our experimental demonstrations are designed to meet this assumption, other methods that identify transitions between contiguous constraints or that incorporate multiple constraints within a single model could be incorporated in the future.

Since measured forces can occur due to other effects including gravitational force or forces induced from accelerating motion, force thresholding is susceptible to false positive segments. To compensate for this we extend force thresholding to include a constraint null hypothesis.

A. Constraint Null Hypothesis

The goal of the constraint null hypothesis is to remove false positive constraint fits. The fitting and model selection procedures described in sections IV and V determine the best model that matches the data for a particular segment. For each of these segments, we evaluate whether the constraint should be kept or discarded (i.e. is it a false positive). An error threshold for each error metric (position, orientation, force and moment) is applied to each sample in the segment and contributes to cumulative tallies V_r , V_q , V_f and V_n . A tally is awarded when the error is more than the threshold. The thresholds are set above the system's measurement noise. If the percentage of cumulative tallies for any metric is greater than an acceptable maximum percentage (e.g. 75% in our implementation), the segment is discarded. Additionally, any segment consisting of only a few samples will easily fit a given model, but does not lend to confidence that the data is actually constrained. For example, if the segment is only two points, the *prismatic relaxed* model will always fit with zero error. To avoid this case, we discard segments of very short duration (in our implementation, less than or equal to five samples).

VII. EXPERIMENTAL EVALUATION AND DISCUSSION

The constraint inference approach is evaluated in three experiments. In the first two experiments, we desire to assess the situations in which wrench information improves constraint fitting of unknown objects. To do so, we compare a method using pose and wrench measurements to an approach that uses only pose measurements. These two experiments include (1)




Fig. 7: Physical constraints used in experiment.

an evaluation of *model fitting* (Section VII-B) for demonstrations containing a single constraint and (2) an evaluation of *model selection* (Section VII-D) for demonstrations containing a single constraint. The third experiment seeks to evaluate the complete constraint inference method on a situated demonstration containing multiple differing constraints.

A. Experimental setup

We desired to show that our methods can reliably recognize and parameterize constraints through a variety of demonstrator strategies. In all three experiments, data was collected from 9 participants with no prior knowledge of or prior experience with the research recruited from the university student population. The study took 60 minutes and all participants received \$12 as compensation. To acquire the resulting pose and wrench measurements that occur during constraint interactions, subjects used custom instrumented tongs [6] equipped with force-torque sensors (ATI mini40) and motion capture markers (Opti-track Flex 13) to perform demonstrations. The instrumented tongs, shown in Fig. 6, provide measured position, orientation, force, and moment trajectories sampled at 120 Hz (after appropriate resampling and filtering). While our method extends to other input methods capable of measuring interaction wrenches, instrumented tongs provides a familiar interface for human demonstrators.

For the first two experiments, where the *model fitting* and *model selection* performance are evaluated, an experimental test bed, incorporating various physical constraints was constructed (see Fig. 7). Each constraint in the test bed has a designated grasp location that participants use to interact with the constraint.

In all demonstrations except *point-on-plane*, *point on line*, and the situated demonstration, users interact with a non-rigid spherical interface. For constraint models that do not incorporate orientation constraints, such as *prismatic relaxed*, *planar relaxed*, and *axial relaxed* constraint, the measured pose information is taken directly from the instrumented tong motion capture markers. For constraint models where orientation of the constrained body is enforced, such as *point-on-plane*, *point-on-line*, *prismatic*, *axial*, and *planar* constraints, the pose information is taken from motion capture makers attached directly to the manipulated object. While our method can fit rigid interactions using the pose of the instrumented

tongs, users would allow rotation when interfacing with the spherical interface, and using separate markers attached to the object allowed for rigid and relaxed data to be captured from a single demonstration without having to instruct users to perform a separate set of demonstrations where objects were held rigidly. In cases where the object markers are used, the applied forces and moments measured using the instrumented tongs are transformed to the object coordinate frame by relating the object and instrumented tongs motion capture markers.

B. Evaluation of Model Fitting

The first experiment seeks to evaluate the performance of constraint *model fitting* (Section IV) for demonstrations containing a single constraint. We compare the performance of a method using pose and wrench information to our previous method [14], [15], which uses only pose measurements.

Specifically, we evaluate fit accuracy as a function of demonstration duration. In the experiments, participants were instructed to interact with each physical constraint on the test bed. Two 10 second demonstrations were recorded per constraint. To simulate varying lengths of constraint interactions, contiguous segments are randomly chosen from across the entire demonstration and fit using the method described in IV. For comparison, the data was also fit to constraint models using a method that considered pose information only. In this case, the method is identical to that described in Section IV except that the maximum likelihood estimate objective function relies on the kinematic error objective given in equation (19) only (i.e., does not include the wrench error objective function from equation (23)).

To compare the methods, a fit error based on the position error metric described in equation (26) is computed for both methods. While a variety of error metrics are possible, we chose a position-based method to allow comparison across constraint models and as a method that can be related to unexpected applied force through a hybrid controller. The fits are evaluated against either the entire recorded demonstration or for *point-on-line* and *point-on-plane*, an independent demonstration performed by the experimenter (since the fitting requires a significant amount of the original 10 second demonstration).

Fig. 8 shows a regression analysis for constraints with large

number of degrees of freedom. Fit accuracy is evaluated as a function of contiguous samples (sampled at 120 Hz). For models with high degrees of freedom (*point-on-plane* - 5 DOF, *planar relaxed* - 4 DOF, *planar* - 3 DOF, and *point-on-line* - 4 DOF) considerably fewer samples of combined kinematic and wrench data are required to fit models than for kinematic (i.e., pose) data alone.

Wrench measurements provide additional information that compensates for degenerate kinematic information allowing a constraint to be fit with fewer samples. Fig. 9 shows a comparison between fitting the *planar relaxed constraint* model with pose and wrench measurements compared to fitting with only pose measurements for a particular participant's demonstration. Short demonstrations on the *planar relaxed constraint* may consist of primarily linear motion which leads to degenerate planes when fitting with only kinematic data. Motion in a straight line will not completely determine the orientation of the plane as many different planes may fit the same line.

For models with fewer degrees of freedom such as the *prismatic*, *axial rotation*, *axial relaxed*, and *planar* constraints, the wrench method yields similar results to the kinematic-only fitting method. This is expected since high numbers of kinematic constraints lead to many directions of admissible forces and moments which limits the utility in fitting. For example, in the *prismatic* model, there is only a single degree of freedom that is not constrained. However, this is the direction of motion which is also the direction in which friction occurs. Since we assume that forces in this direction could be friction, they are removed. The remaining two directions are acceptable reaction forces. As a result, there is no additional information added from including wrench measurements in fitting for this model. Fig. 10 shows a regression analysis of number of samples verses fit error for the constraints with a small number of degrees of freedom. The *axial relaxed* and *axial rotation* models have high nonlinearity and sometimes the fits are incorrect due to convergence to a local minimum. This explains the large standard deviation for both models even when using a large number of samples.

In this experimental setup, the orientation of the force frame is derived from motion capture orientation measurements. As a result, the force and moment errors tend to be higher based on the propagated orientation error. This can be seen in the increased noise floor for the wrench fitting method in the regression plots (i.e., the wrench method converges to a higher error than the kinematic method). While this is a product of the experimental setup, it demonstrates that it may still be beneficial to consider using kinematic fitting for constraints where the force and moments add little information (e.g., *prismatic constraint*) to avoid cases when the forces and moments have a significant noise floor and to simplify the dimensionality of the non-linear optimization.

The experiments show that wrench measurements improve fitting of models with high degrees of freedom over short demonstration duration. They help in situations when kinematic information is insufficient to completely define the constraint model. For constraint models with high degrees of freedom, we recommend that wrench measurements are




Fig. 8: Comparison between fit accuracy in meters versus number of contiguous samples (120 Hz) in a demonstration for high degree of freedom models. Fitting with kinematic and wrench measurements [green line]. Fitting with kinematic measurements [red line]. Shaded regions correspond to the first standard deviation about the mean. Constraints with higher degrees of freedom converge with fewer samples when wrench measurements are incorporated during the fitting process.

incorporated in the fitting process. For models with low degrees of freedom, wrench measurements can be omitted to decrease computation time and reduce the complexity of the optimization process.

C. Evaluation of Model Selection

The second experiment seeks to evaluate the performance of constraint *model selection* (Section V) for demonstrations containing a single constraint and compare the performance of a method using both pose and wrench measurements to an approach that uses only pose measurements. For comparison, we also include a penalized kinematic method [24] to address limitations with a kinematic-only method, where demonstra-




Fig. 9: An example demonstration of the *planar relaxed constraint* from Demonstrator 6. Using a small portion in the demonstration (cyan) of 0.2 seconds fitting with only kinematic data (red plane) does not produce the correct model. Including reaction wrenches in the analysis fits a better plane (green) to the demonstration.

tion data may be consistent with a number of different kinematic models and thus result in poor selection accuracy. The penalized likelihood metric is defined as follows:

$$m = \arg \min_m (\beta \text{DOF}(m) + \gamma \chi - \ln L(m)) \quad (32)$$

where DOF refers to the kinematic degrees of freedom of a model and χ is the number of model parameters (Corresponding values for the constraint models used in this paper are provided in Appendix A. Weighted parameters β and γ are fit using eighty training demonstrations (10 per constraint model). The accuracy is maximized with $\beta = 1.69$ and $\gamma = 0.40$ (accuracy: 0.65).

In our experiments, the demonstrations are split into 3.33 second segments (sufficient samples for fitting). Since the demonstrations are sufficiently long, the constraint parameters for all demonstrations are fit with kinematic data to ensure that the differing methods use the same constraint fit parameters. Since model selection is requisite on the constraint being properly fit, fits that do not match the performed constraint are manually discarded. The results appear in Fig. 11.

The results shown in Fig. 11 were, in part, evaluated using the error tallying method described in equation (30) which, in turn, requires that we set the thresholds for each of the error metrics. The metric values are set based on an estimate of the data collection accuracy. The force and moment error represent the component orthogonal to the permissible force and moment directions. Since the errors are normalized with respect to the measured force and moment, the formulation is equivalent to evaluating the sine of the angle between the measured and normal vectors. In our experimental evaluation, the error for the force and moment is set to the same value as the orientation error (2 degrees). The list of parameter values is provided in Table I below.

Threshold	Value
Position	2 mm
Orientation	2 deg
Force	0.06 (sin(2 deg))
Moment	0.06 (sin(2 deg))

TABLE I: Threshold values for error tallying




Fig. 10: Fit accuracy in meters versus number of contiguous samples (120 Hz) in a demonstration for each constraint. Fitting with kinematic and wrench measurements [green line]. Fitting with kinematic measurements [red line]. Shaded regions correspond to the first standard deviation about the mean. Constraints with lower degrees of freedom do not benefit when wrench measurements are incorporated during the fitting process.

Our results demonstrate the value of including wrenches in constraint model selection. As seen in Figure 11, the method including wrench information in model selection greatly outperforms equivalent kinematic and penalized-kinematic voting methods. In the kinematic-only tallying method, incorrect classifications can be attributed to model generality. For example, the *axial relaxed* constraint is always identified as *planar relaxed*. When only kinematic measurements are considered, this is a likely case. There is always a plane that can describe an axial motion and moreover any small errors in the position might still be consistent with the plane model, but not necessarily with the axial model. Two other models, (*planar constraint* and *prismatic constraint*), are consistently

constrained motion. Additionally, the espresso machine lever (*axial relaxed constraint*) can be difficult to grasp leading to some demonstrations that do not resemble axial motion. As expected, a few constraint interactions were misclassified by model selection or rejected by the null hypothesis. In spite of the imprecise constraints, the system correctly infers a majority of the constraints.

The null hypothesis was successful in rejecting non-constrained motion that exceeded the force threshold, such as when participants hold the coffee pod and accelerate towards the espresso maker. In the 18 demonstrations, the null hypothesis rejected 20 extraneous segments.

Fig. 1 shows the results for one of the participants where the method is successful. Our approach correctly selected the model 77% of the time. Correctly inferred constraints were all correctly fit (as shown in Fig 1). Constraints were identified using a single demonstration in which the demonstrator performed them naturally with minimal instruction.

VIII. CONCLUSION

The limitations of our method suggest future areas of extension. In our methods, we assume free space interaction occurs between each of the constrained segments. While we have not explored how often this assumption is valid, other work has developed methods for detecting changes between consecutive constraints.

In this work, we focus on methods for constraint inference, however, we do not explore task replication by a robot. While our method provides a general set of tools to recognize constraints and their associated parameters, it is unanswered whether these models appropriately define a task such that a robot could perform similar actions.

Additionally, this work does not discuss human procedures that may be identified and leveraged during replay. We also do not explore remapping of constrained motions between a general input method and a robot which may have different kinematic restrictions. Our method describes when constraints occur and their parameters for a given demonstration. Future work will investigate editing these representations, for example to modify a motion on a constraint.

In summary, we present a framework that is capable of recognizing geometric constraints in a demonstration. Specifically, the method includes the ability to segment sections of constrained motion, fit models using both kinematic and wrench measurements, and select the best candidate model for each section. We show that wrench measurements enable estimation of constraints with high degrees of freedom from fewer samples compared to using only kinematic measurements. We demonstrate the use of wrench measurements in improving model selection classification accuracy compared to a similar method using only kinematic measurements. We demonstrate that model selection can be further improved by incorporating model generality through the degrees of freedom and number of parameters of candidate constraint models. We evaluate our end-to-end constraint inference method on a situated task of preparing an espresso which correctly identifies 77% of constraints without any user input.

APPENDIX A

MATHEMATICAL MODELS OF GEOMETRIC CONSTRAINTS

This section describes the mathematical model development for the various constraint types evaluated in the paper (see Fig. 3). The mathematical model for the *point-on-plane constraint* is provided in section III-A1.

For all constraint types, our mathematical model considers a 6-degree of freedom rigid body of negligible inertial properties located in space through 3 translational coordinates $\mathbf{r} \in \mathbb{R}^3$ and rotational coordinates represented either by a unit quaternion $\mathbf{q} \in \text{Spin}(3)$ or an orthogonal rotational matrix $\mathbf{A}(\mathbf{q}) \in \text{SO}(3)$. Together $(\mathbf{r}, \mathbf{q})$ form the body coordinates $\mathbf{p}$. Each constraint type has a set of parameters, $\boldsymbol{\alpha}$, that parameterize the geometry of the constraint.

A. Planar constraint

The planar constraint is exemplified by an eraser moving against a whiteboard. The rigid body interacting with the plane (e.g. eraser) can only rotate about a vector perpendicular to the plane, and all points within the rigid body translate parallel to this plane.

A plane may be specified by applying a general displacement (i.e. translation and rotation) transformation of the x - y plane which involves:

- 1) translating the x - y coordinate plane along the z -axis
- 2) rotating the translated plane about the origin.

We represent the rotation transformation using two exponential coordinates $\mathbf{w} = [w_x, w_y, 0] \in \mathbb{R}^2$ corresponding to $e^{\tilde{\mathbf{w}}} \in \text{SO}(3)$, the exponential map, which is equivalent to a rotation matrix with an axis of rotation in the x - y plane. Rodrigues' rotation formula [17] is used to compute $e^{\tilde{\mathbf{w}}}$. $\tilde{\mathbf{w}}$ is the skew symmetric matrix similar to equation (62) in Appendix B. The third term of $\mathbf{w}$ is zero because rotations about the z - axis (perpendicular to the plane) do not alter the plane's geometry. The translation is represented by $d \in \mathbb{R}$. The normal vector on this plane is represented by $e^{\tilde{\mathbf{w}}}[0 \ 0 \ 1]^T$ and the shifted origin of the x - y plane is represented by $e^{\tilde{\mathbf{w}}}[0 \ 0 \ d]^T$.

The origin of the local coordinate frame on the body is coincident with the plane. The constraint equations for motion on a plane is given as:

$$\Phi_1 \equiv \left(e^{\tilde{\mathbf{w}}}[0 \ 0 \ d]^T - \mathbf{r} \right)^T e^{\tilde{\mathbf{w}}}[0 \ 0 \ 1]^T = 0 \quad (33)$$

A unit vector $\bar{\mathbf{t}}$ is defined to force the rigid body perpendicular to the plane:

$$\Phi_2 \equiv (\mathbf{A}(\mathbf{q})\bar{\mathbf{t}})^T e^{\tilde{\mathbf{w}}}[1 \ 0 \ 0]^T = 0 \quad (34)$$

$$\Phi_3 \equiv (\mathbf{A}(\mathbf{q})\bar{\mathbf{t}})^T e^{\tilde{\mathbf{w}}}[0 \ 1 \ 0]^T = 0 \quad (35)$$

$$\Phi_{p1} \equiv \bar{\mathbf{t}}^T \bar{\mathbf{t}} - 1 = 0 \quad (36)$$

Unity of $\bar{\mathbf{t}}$ is enforced by equation (36). The parameters of this constraint are $\boldsymbol{\alpha} = (\bar{\mathbf{t}}, d, w_x, w_y)$. The rigid body has 3 degrees of freedom.

Friction Model: The friction force is in the direction of linear velocity and the friction moment is in the direction of angular velocity.

$$\mathbf{F}_\mu = (\mathbf{F} \cdot \hat{\mathbf{v}})\hat{\mathbf{v}} \quad (37)$$

$$\mathbf{N}_\mu = (\mathbf{N} \cdot \hat{\boldsymbol{\omega}})\hat{\boldsymbol{\omega}} \quad (38)$$

where $\mathbf{F}$ and $\mathbf{N}$ are the measured force and moment, respectively. The measured linear and angular velocity are $\mathbf{v}$ and $\boldsymbol{\omega}$, respectively.

B. Planar relaxed constraint

The *planar relaxed constraint* is represented by equation (33). The friction model is the same as the *planar constraint*.

C. Prismatic constraint

The prismatic constraint represents translational motion in one direction. It is similar to pulling out a drawer. All points on the rigid body translate identically. This constraint requires a representation of the axis of translation. Similar to the plane, the axis of translation can be generated by applying a general displacement (i.e. translation and rotation) transformation of the z coordinate axis which is equivalent to:

- 1) translating the z axis in the x - y plane
- 2) rotating the translated axis about the origin.

This is represented by two exponential coordinates $\mathbf{w} = [w_x, w_y, 0] \in \mathbb{R}^2$ and two translational coordinates $d_x, d_y \in \mathbb{R}$. The third term in $\mathbf{w}$ is zero because rotations about the z -axis produce a line that could be produced by an alternative translation motion. This defines the axis in the global reference frame. The tangent to this axis is $e^{\tilde{\mathbf{w}}}[0 \ 0 \ 1]^T$ and the translated origin is $e^{\tilde{\mathbf{w}}}[d_x \ d_y \ 0]^T$. $e^{\tilde{\mathbf{w}}}[1 \ 0 \ 0]^T$ and $e^{\tilde{\mathbf{w}}}[0 \ 1 \ 0]^T$ represent vectors perpendicular to this axis. The origin of the local coordinate frame of the body is coincident with this line/axis:

$$\Phi_1 \equiv \left(e^{\tilde{\mathbf{w}}}[d_x \ d_y \ 0]^T - \mathbf{r} \right)^T e^{\tilde{\mathbf{w}}}[1 \ 0 \ 0]^T = 0 \quad (39)$$

$$\Phi_2 \equiv \left(e^{\tilde{\mathbf{w}}}[d_x \ d_y \ 0]^T - \mathbf{r} \right)^T e^{\tilde{\mathbf{w}}}[0 \ 1 \ 0]^T = 0 \quad (40)$$

The orientation of the rigid body is fixed using the following equations:

$$\Phi_3 \equiv (\mathbf{A}(\mathbf{q})\bar{\mathbf{s}})^T e^{\tilde{\mathbf{w}}}[1 \ 0 \ 0]^T = 0 \quad (41)$$

$$\Phi_4 \equiv (\mathbf{A}(\mathbf{q})\bar{\mathbf{s}})^T e^{\tilde{\mathbf{w}}}[0 \ 1 \ 0]^T = 0 \quad (42)$$

$$\Phi_5 \equiv (\mathbf{A}(\mathbf{q})\bar{\mathbf{t}})^T e^{\tilde{\mathbf{w}}}[1 \ 0 \ 0]^T = 0 \quad (43)$$

$$\Phi_{p1} \equiv \bar{\mathbf{s}}^T \bar{\mathbf{s}} - 1 = 0 \quad (44)$$

$$\Phi_{p2} \equiv \bar{\mathbf{s}}^T \bar{\mathbf{t}} = 0 \quad (45)$$

where $\bar{\mathbf{s}}$ and $\bar{\mathbf{t}}$ are unit vectors fixed on the rigid body. Equations (41) through (36) and (44) prevent the body from rotating. Equations (44) and (36) forces unity of $\bar{\mathbf{s}}$ and $\bar{\mathbf{t}}$ respectively. This rigid body has one degree of freedom of motion. The constraint parameters are: $\boldsymbol{\alpha} = (\bar{\mathbf{t}}, \bar{\mathbf{s}}, d_x, d_y, w_x, w_y)$. **Friction Model:** Friction force is in the direction of linear velocity. This model does not permit angular rotation so we assume zero friction moment.

$$\mathbf{F}_\mu = (\mathbf{F} \cdot \hat{\mathbf{v}})\hat{\mathbf{v}} \quad (46)$$

$$\mathbf{N}_\mu = 0 \quad (47)$$

where $\mathbf{F}$ is the measured force and $\mathbf{v}$ is the linear velocity.

D. Prismatic relaxed constraint

The *prismatic relaxed constraint* is obtained by only using equations (39) and (40). The friction model is given by equations (46) and (38).

E. Point-on-line constraint

The *point-on-line constraint* represents a point on a rigid body constrained to a line in space. It is similar to a pen moving against a straight edge.

Consider a point $\mathbf{s}^* \in \mathbb{R}^3$ in the global reference frame defined as a fixed point on the rigid body. $\mathbf{s}$ is a vector directed from the origin of the local reference frame of the body to the rigidly fixed point $\mathbf{s}^*$ in the global reference frame. The local reference frame counterpart of $\mathbf{s}$ is $\bar{\mathbf{s}}$:

$$\mathbf{s}^* \equiv \mathbf{r} + \mathbf{s} \equiv \mathbf{r} + \mathbf{A}(\mathbf{q})\bar{\mathbf{s}} \quad (48)$$

where $\mathbf{A}(\mathbf{q})$ is the orthogonal rotation matrix equivalent to the rotation represented by $\mathbf{q}$. $\mathbf{s}^*$ on the rigid body is coincident with a line represented by parameters (d_x, d_y, w_x, w_y) :

$$\Phi_1 \equiv \left(e^{\tilde{\mathbf{w}}}[d_x \ d_y \ 0]^T - \mathbf{r} - \mathbf{A}(\mathbf{q})\bar{\mathbf{s}} \right)^T e^{\tilde{\mathbf{w}}}[1 \ 0 \ 0]^T = 0 \quad (49)$$

$$\Phi_2 \equiv \left(e^{\tilde{\mathbf{w}}}[d_x \ d_y \ 0]^T - \mathbf{r} - \mathbf{A}(\mathbf{q})\bar{\mathbf{s}} \right)^T e^{\tilde{\mathbf{w}}}[0 \ 1 \ 0]^T = 0 \quad (50)$$

Friction Model: The friction models are identical to the friction models for the *point-on-plane constraint* described in section III-A1.

F. Axial rotation constraint

The axial rotation constraint is similar to a door knob or a hinged door, all points on the rigid body rotate about an axis and translations are not permitted.

Consider a point on the rigid body $\mathbf{s}^*$ that is on the axis of rotation and in the plane perpendicular to the axis containing the coordinate frame origin. A rigid point on the axis of rotation defined as $(d_x, d_y, d_z) \in \mathbb{R}^3$ constrains the point $\mathbf{s}^*$. The vector $\bar{\mathbf{s}}$ is perpendicular to this axis. The rigid body must not rotate about vector $\bar{\mathbf{s}}$ so unit vector $\bar{\mathbf{t}}$ is introduced to enforce this condition. The constraint equations are (36), (45) and the following:

$$\Phi_1 \equiv \mathbf{r} + \mathbf{A}(\mathbf{q})\bar{\mathbf{s}} - [d_x, d_y, d_z]^T = 0 \quad (51)$$

$$\Phi_2 \equiv (\mathbf{A}(\mathbf{q})\bar{\mathbf{t}})^T e^{\tilde{\mathbf{w}}}[0 \ 0 \ 1]^T = 0 \quad (52)$$

$$\Phi_3 \equiv (\mathbf{A}(\mathbf{q})\bar{\mathbf{s}})^T e^{\tilde{\mathbf{w}}}[0 \ 0 \ 1]^T = 0 \quad (53)$$

Constraint parameters are $\boldsymbol{\alpha} = (\bar{\mathbf{t}}, \bar{\mathbf{s}}, d_x, d_y, d_z, w_x, w_y)$. This constraint has 1 degree of freedom.

Friction Model: The friction wrenches and reaction wrenches are coupled in this model. The principle of virtual work enables the computation of reaction forces and reaction wrenches from the corresponding measured quantities. Let's assume all measured forces are reaction forces. This implies that friction forces are zero:

$$\mathbf{F}_\mu = 0 \quad (54)$$

The principle of virtual work requires that wrenches are orthogonal to twists as per Equation 7 (substituting virtual displacements with velocities). The residual of equation 7 provides the corresponding friction work W_μ :

$$W_\mu = \mathbf{v} \cdot \mathbf{F} + \boldsymbol{\omega} \cdot \mathbf{N} \quad (55)$$

The friction work comes from the friction moment about the axis of rotation. The friction moment is directed about the axis of rotation.

$$\mathbf{N}_\mu = \frac{W_\mu}{|\boldsymbol{\omega}|} \hat{\boldsymbol{\omega}} \quad (56)$$

G. Axial relaxed constraint

The *axial relaxed constraint* also has parameters $(d_x, d_y, d_z, w_x, w_y)$ to configure the axis of rotation and the center of rotation. The radius $l \in \mathbb{R}$ is included to fully define the constraint.

$$\Phi_2 \equiv \left\| \mathbf{r} - [d_x, d_y, d_z]^T \right\| - l^2 = 0 \quad (57)$$

$$\Phi_1 \equiv \left(\mathbf{r} - [d_x, d_y, d_z]^T \right)^T e^{\tilde{\mathbf{w}}} [0 \ 0 \ 1]^T = 0 \quad (58)$$

Equation (57) restricts the body origin to a sphere of radius l centered about (d_x, d_y, d_z) . Equation (58) restricts the body origin to a plane containing (d_x, d_y, d_z) with normal $e^{\tilde{\mathbf{w}}} [0 \ 0 \ 1]^T$.

The friction model is represented by equations (37) and (38)

H. Degrees of freedom and model parameters per constraint

Table II shows the degrees of freedom and number of constraint parameters for each of the models considered in this paper. These values are used to construct the penalized likelihood method used for comparison in section VII-C.

Constraint	DOF	number of model parameters
Point-on-plane	5	6
Point-on-line	4	7
Planar	3	6
Planar relaxed	5	3
Prismatic	1	10
Prismatic relaxed	4	4
Axial rotation	1	11
Axial relaxed	4	8

TABLE II: Degrees of freedom (DOF) and number of model parameters for each constraint

APPENDIX B

RELATIONSHIP BETWEEN π AND $\mathbf{q}$

Evaluating the partial derivatives of the constraint equations, Φ , is required to determine admissible reaction forces and moments for a given constraint. The partial derivative with respect to rotation uses the virtual rotation variable, $\boldsymbol{\pi}$. This section describes the derivation of $\boldsymbol{\pi}$ in terms of the known orientation, $\mathbf{q}$.

$$\Phi_{\pi} = \Phi_{\mathbf{q}} \frac{1}{2} \mathbf{G}^T \mathbf{A}^T \quad (59)$$

where the rotation is described by quaternion $\mathbf{q} = [e_0 \ \mathbf{e}^T]^T$ and where $e_0 = q_0$ and $\mathbf{e} = [q_1 \ q_2 \ q_3]^T$. $\mathbf{A}$ is the orthogonal rotation matrix and $\mathbf{G}$ is a matrix constant.

$$\mathbf{A} =: (e_0^2 - \mathbf{e}^T \mathbf{e}) \mathbf{I} + 2\mathbf{e}\mathbf{e}^T + 2e_0\tilde{\mathbf{e}} \quad (60)$$

$$\mathbf{G} =: \begin{bmatrix} -\mathbf{e} & -\tilde{\mathbf{e}} + e_0 \mathbf{I} \end{bmatrix} \quad (61)$$

where

$$\tilde{\mathbf{e}} = \begin{bmatrix} 0 & -e_z & e_y \\ e_z & 0 & -e_x \\ -e_y & e_x & 0 \end{bmatrix} \quad (62)$$

and $\mathbf{I}$ is the identity matrix

ACKNOWLEDGMENT

This work was supported in part by NSF award 1830242 and the University of Wisconsin- Madison Office of the Vice Chancellor for Research and Graduate Education with funding from the Wisconsin Alumni Research Foundation.

REFERENCES

- [1] D. Berenson, S. S. Srinivasa, D. Ferguson, and J. J. Kuffner, "Manipulation planning on constraint manifolds," *2009 IEEE International Conference on Robotics and Automation*, vol. i, pp. 625–632, 2009. [Online]. Available: <https://ieeexplore.ieee.org/document/5152399>
- [2] M. Stilman, "Task constrained motion planning in robot joint space," *IEEE International Conference on Intelligent Robots and Systems*, pp. 3074–3081, 2007.
- [3] J. Huh, B. Lee, and D. D. Lee, "Constrained Sampling-Based Planning for Grasping and Manipulation," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 5 2018, pp. 223–230. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8461265>
- [4] M. H. Raibert and J. J. Craig, "Hybrid Position/Force Control of Manipulators," *Journal of Dynamic Systems, Measurement, and Control*, vol. 103, no. 2, p. 126, 1981.
- [5] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469–483, 2009. [Online]. Available: <https://doi.org/10.1016/j.robot.2008.10.024>
- [6] P. Praveena, G. Subramani, B. Mutlu, and M. Gleicher, "Characterizing Input Methods for Human-to-robot Demonstrations," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 3 2019, pp. 344–353. [Online]. Available: <http://graphics.cs.wisc.edu/Papers/2019/PSMG19>
- [7] L. Armesto, J. Moura, V. Ivan, M. S. Erden, A. Sala, and S. Vijayakumar, "Constraint-aware learning of policies by demonstration," *The International Journal of Robotics Research*, vol. 37, no. 13-14, pp. 1673–1689, 2018. [Online]. Available: <https://doi.org/10.1177/0278364918784354>
- [8] M. Li, K. Tahara, and A. Billard, "Learning task manifolds for constrained object manipulation," *Autonomous Robots*, vol. 42, no. 1, pp. 1–16, 2017.
- [9] A. Rodríguez, L. Basañez, and E. Celaya, "A Relational Positioning Methodology for Robot Task Specification and Execution," *IEEE Transactions on Robotics*, vol. 24, no. 3, pp. 600–611, 6 2008.
- [10] S. Dutre, H. Bruyninckx, and J. De Schutter, "Contact identification and monitoring based on energy," in *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2. IEEE, 1996, pp. 1333–1338.
- [11] J. de Schutter, H. Bruyninckx, S. Dutré, J. de Geeter, J. Katupitiya, S. Demey, and T. Lefebvre, "Estimating First-Order Geometric Parameters and Monitoring Contact Transitions during Force-Controlled Compliant Motion," *The International Journal of Robotics Research*, vol. 18, no. 12, pp. 1161–1184, 1999. [Online]. Available: <https://doi.org/10.1177/02783649922067780>
- [12] T. Lefebvre, H. Bruyninckx, and J. Schutter, "Polyhedral Contact Formation Modeling and Identification for Autonomous Compliant Motion," *Robotics and Automation, IEEE Transactions on*, vol. 19, pp. 26 – 41, 2 2003.
- [13] W. Meeussen, J. Rutgeerts, K. Gadeyne, H. Bruyninckx, and J. De Schutter, "Contact state segmentation using particle filters for programming by human demonstration in compliant motion tasks," *Springer Tracts in Advanced Robotics*, vol. 39, pp. 3–12, 2008.
- [14] G. Subramani, M. Zinn, and M. Gleicher, "Inferring geometric constraints in human demonstrations," *Proceedings of the 2nd Annual Conference on Robot Learning*, no. CoRL, pp. 1–14, 2018. [Online]. Available: <http://arxiv.org/abs/1810.00140>
- [15] G. Subramani, M. Gleicher, and M. Zinn, "Recognizing Geometric Constraints in Human Demonstrations using Force and Position Signals," *IEEE Robotics and Automation Letters*, vol. 3766, no. c, pp. 1–1, 2018. [Online]. Available: <http://ieeexplore.ieee.org/document/8264714/>
- [16] H. Goldstein, C. P. Poole, and J. L. Safko, *Classical Mechanics*. Addison Wesley, 2002. [Online]. Available: <https://books.google.com/books?id=uCuQgAACAAJ>
- [17] R. M. Murray, Z. Li, S. S. Sastry, and S. S. Sastry, *A mathematical introduction to robotic manipulation*. CRC press, 1994.
- [18] E. J. Haug, *Computer aided kinematics and dynamics of mechanical systems*. Allyn and Bacon Boston, 1989, vol. 1.

- [19] H. M. Choset, S. Hutchinson, K. M. Lynch, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun, *Principles of robot motion: theory, algorithms, and implementation*. MIT press, 2005.
- [20] C. Lanczos, *The variational principles of mechanics*. Courier Corporation, 2012.
- [21] C. S. Burrus, "Iterative reweighted least squares," *OpenStax CNX*. Available online: <http://cnx.org/contents/92b90377-2b34-49e4-b26f-7fe572db78a1>, vol. 12, 2012.
- [22] R. Fletcher, *Practical methods of optimization*. John Wiley & Sons, 2013.
- [23] S. Niekum, S. Osentoski, C. G. Atkeson, and A. G. Barto, "Online Bayesian changepoint detection for articulated motion models," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2015-June, no. June, pp. 1468–1475, 2015.
- [24] H. Akaike, "A new look at the statistical model identification," *IEEE Transactions on Automatic Control*, vol. 19, no. 6, pp. 716–723, 12 1974.



Michael Zinn (M'07) received the B.S. and M.S. degrees from Massachusetts Institute of Technology, Cambridge, MA, USA, in 1987 and 1988, respectively, and the Ph.D. degree in mechanical engineering from Stanford University, Stanford, CA, USA, in 2005. He joined the faculty at the University of Wisconsin, Madison, WI, USA, in 2007. Prior to joining the University of Wisconsin, Madison faculty, he was the Director of Systems and Controls Engineering at Hansen Medical where he helped to develop the world's first commercially available minimally invasive flexible surgical robotic system. In addition to his experience at Hansen Medical, he has more than ten years of electro-mechanical system design and manufacturing experience in aerospace and high-technology industries. His research interests include understanding and overcoming the design and control challenges of complex electro-mechanical systems with a primary focus on human-centered robotics. His focus on human-centered robotics spans multiple application areas including manufacturing, medical devices, and haptics.



Guru Subramani received the B.S. degree in mechanical engineering from National Institute of Technology Tiruchirappalli in 2012 and the M.S. and Ph.D. degrees in mechanical engineering from University of Wisconsin-Madison in 2014 and 2019, respectively.



Michael Hagenow received the B.S. degree in mechanical engineering from Tufts University in 2014 and the M.S. degree in mechanical engineering from University of Wisconsin-Madison in 2019, where he is working towards the Ph.D. degree in mechanical engineering. His research interests include robotics, control, and dynamics.



Michael Gleicher is a Professor in the Department of Computer Sciences at the University of Wisconsin, Madison. Prof. Gleicher is founder of the Department's Visual Computing Group. His research interests span the range of visual computing, including data visualization, robotics, image and video processing tools, virtual reality, and character animation. His current foci are human data interaction and human robot interaction. Prior to joining the university, Prof. Gleicher was a researcher at The Autodesk Vision Technology Center and in Apple

Computer's Advanced Technology Group. He earned his Ph. D. in Computer Science from Carnegie Mellon University, and holds a B.S.E. in Electrical Engineering from Duke University. In 2013-2014, he was a visiting researcher at INRIA Rhone-Alpes. Prof. Gleicher is an ACM Distinguished Scientist.