Proceedings of the 2012 IEEE
International Conference on Robotics and Biomimetics
December 11-14, 2012, Guangzhou, China

Gradient Calibration for the RCBHT Cantilever Snap Verification System.

J. Rojas, K. Harada, H. Onda, N. Yamanobe, E. Yoshida, K. Nagata, and Y. Kawai
Intelligent Sys. Research Institute, AIST
Tsukuba, Ibaraki, 305-8568, Japan

Abstract— **In this work a gradient calibration method was presented as part of the Relative-Change-Based-Hierarchical Taxonomy (RCBHT) cantilever-snap verification system and the Pivot Approach control strategy for the automation of cantilever-snaps. As part of a relative-change based force signal interpretation scheme, an effective gradient calibration process is needed to increase the RCBHT's system robustness. Prior to this work, all gradient classification schemes were derived on an intuitive trial and error basis. Statistical measures were used to derive contact and constant gradient thresholds in contextually sensitive ways. The method requires training assemblies to identify a minimum contact gradient which serves as a marker for all other gradient thresholds. Experimental procedures verified that our calibration method was effective. Assemblies with supervised successful outcomes were used in experimentation. The RCBHT assessed out assemblies as successful using the calibration method. Even two snaps that where classified falsely as unsuccessful when using a previously non-calibrated version of the RCBHT.**

I. Introduction

As part of our ongoing efforts to automate cantilever-snap assemblies we have implemented a number of features consisting of: (a) the Pivot Approach (PA) control strategy for cantilever-snap type assemblies. The PA exploits snap parts' hardware design to constraint the task's motion and generate similar sensory-signal patterns across trials and systematically discretize the assembly into intuitive states (see Sec. III); (b) the Relative-Change Based Hierarchical Taxonomy (RCBHT) system for snap verification [1]. The latter worked in concert with the PA and was built on the premise that relative-change patterns can be classified through a small category set while aided by contextual information (see Sec. IV). The latter has been used to effectively assess the outcome of cantilever-snap assemblies with various degrees of complexities; (c) the probabilistic version of the RCBHT, which used a Bayesian filter to yield beliefs at different levels of the system [2]. The probabilistic version of our system yielded more intuitive outcome representations and used a verification scheme to determine if a task was successful or not.

To this date however, the RCBHT system's first layer uses a gradient classification method that has been derived by trial-and-error. When a new robot or a new snap part is used, we have been forced to run assembly trials, study the data manually and determine which gradient thresholds would be most appropriate. The process is not robust and in our last work [2] it led to the presence of false-positive results for the outcome assessment. In this work we studied whether a contextually sensitive gradient calibration method could be used such that a set of thresholds would work under the RCBHT to properly assess the outcome of successful assemblies across trials.

The calibration method was derived by using statistical measures to derive contact and constant gradient thresholds in contextually sensitive ways. During our experimentation we discovered that by running a number of training trials and selecting the gradient thresholds that belong to the trial whose contact threshold is the lowest, then the gradient calibration would be more effective. The experimentation also revealed a hidden fact prior to this work. That is, that when using three rounds of filtering at the motion composition level, some key motion composition actions are filtered away by the system. Two clean up cycles proved to be a very effective number for correctly assessing snap assembly outcomes. It is also worth mentioning, that the calibration method can work effectively when selecting a different value for 'determination coefficient' described in Sec. IV-A so as long as the same coefficient is used throughout all snap assemblies.

The gradient calibration method was effectively used by the RCBHT system to correctly assess the outcome of 100% training and test snap-assemblies executed in this work. The gradient calibration method provides an important improvement to the previous trial-and-error determination of gradient thresholds for the RCBHT system.

The paper is organized as follows: in Sec. II the experimental setup is described. In Sec. III the PA control strategy is introduced. In Sec. IV the snap verification system is presented. In Sec. V the gradient calibration approach is described. In Sec. V-B experimental procedures are outlined along with results. In Sec. VI important contributions and limitations are discussed, and in Sec. VII key points are summarized.

II. Experimental Setup

In this work the Pivot Approach [3] and the RCBHT system [1] were applied to a dual-arm 6 DoF anthropomorph HIRO robot that was simulated using the OpenHRP environment [4]. A CAD derived camera part consisting of both male and female parts was used. The female part was rigidly held by a specially designed tool mounted on the robot's wrist, while the male mold part, which consisted of four cantilever-snaps (two snaps were used in our previous work)

978-1-4673-2127-3/12/$31.00 © 2012 IEEE

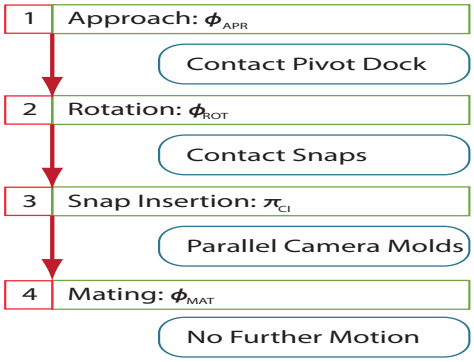

Fig. 2. Each of the four automata states (red) is accompanies by a controller template (green). Transition conditions are specified in blue.

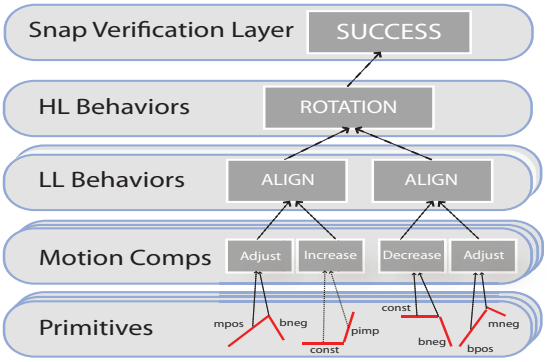

Fig. 3. The RCBHT abstracts FT sensor data to produce intuitive higher-level behavior representations that enable reasoning about the snap assembly's state.

was rigidly fixed to the ground. Cantilever-snap assemblies were executed through the PA and its FT signals interpreted through the RCBHT.

III. Pivot Approach

The PA exploits snap parts' hardware design to constraint the task's motion and generate similar sensory-signal patterns across trials and systematically discretize the assembly into intuitive states [3].

The PA decomposed the assembly into four intuitive automata states as seen in Fig. 2 and used the action transitions seen in Fig. 1. A detailed description of the Pivot Approach is found in [3].

IV. Relative Change-Based Hierarchical Taxonomy

The RCBHT is a state estimation scheme for used in snap assemblies. Over the last two decades much work has been done in active sensing for compliant motion tasks; initially most work was applied to peg-in-hole tasks [5], [6]. Recently, the concept of contact-state graphs has been used for general compliant motion tasks for simple geometrical parts as in [7], [8]. Even with simple geometrical parts, the contact state number can explode. The approach becomes unfeasible for geometrically complex parts as is typically the case with snap contacts.

With this in mind, the RCBHT yields state representations by hierarchically abstracting snap assembly FT data to generate intuitive HLBs [1]. The hierarchical taxonomy is composed of five increasingly abstract layers that encode relative-change in the task's force signatures. The taxonomy is built on the premise that relative-change patterns can be classified through a small set of categoric labels and aided by contextual information. The RCBHT analyzes FT signatures from all force axes independently and contextualizes the state according to automata state participation (the Approach stage is not considered as no FT data is gathered there).

In this section, we will describe the five different layers of the RCBHT but will place more emphasis on the first layer which is the one in which the gradient classifications take place.

A. Primitive Layer

The primitives layer requires that each signal is partitioned into linear segments of data that closely approximate the original signal. Linear regression in concert with a correlation measure (the determination coefficient R^2) is used to partition the data whenever a minimum correlation threshold is crossed. If the determination coefficient drops under a given threshold the linear fit is partitioned and a new regression is started. The R^2 coefficient is a correlation measure that studies the ratio of the sum of the squares of the residual errors between the original data y and the fit data $\hat{y}$ to the sum of the variance ${\sigma_y}^2$ as shown in Eqtn. 1.

$$R^2 = 1 - \frac{\sum (y - \hat{y})^2}{{\sigma_y}^2} \tag{1}$$

In our previous work, we set the threshold used to partition the data was set at 0.70, such that if the correlation dropped to under 70%, a linear segment or "partition" would be generated, and a new one would start at the next data point. The data was traversed by a window equal to five data points (the data was sampled at a frequency of 1kHz by the simulation). The threshold values and the window length were empirically selected to partition the data sufficiently to capture relevant changes in the signals.

Each partition was accompanied by a data structure with seven types of information about itself: the average value across data points, the maximum value, the minimum value, the start time, the end time, the gradient value, and a gradient label. With respect to the latter, nine gradient labels (positive impulse, 'pimp'; big, medium, and small positive gradients, 'b/m/spos'; constant gradients, 'const'; and their negative equivalents, 'nimp', 'b/m/s/neg') were assigned according to ranges summarized in table of Fig. 4. In this work, we defined a gradient calibration method that is described in Sec. V. The classification first attempts to separate instances of data in which contact or mating takes place. On the one hand, contact phenomena is characterized by very rapid and large changes in force signals, almost approximating an impulse. To this end, positive and negative impulses were categorized for gradients with values greater or less than 70. On the other

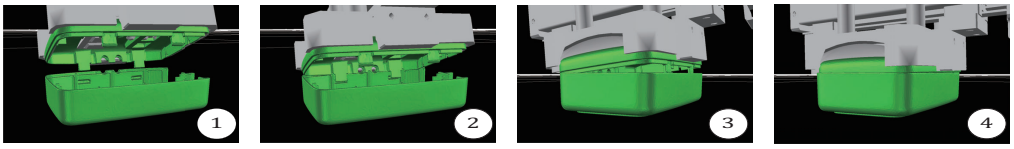

Fig. 1. The Pivot Approach is composed of four states: Approach, Rotation, Snap, and Mating.

hand, for mating and other situations, there is little or no change in force, for this reason a constant label was assigned to signals with gradient values less than the absolute value of 2. In between these two extremes we chose to have three gradient categories for both positive and negative signals to give a general idea of the magnitude change registered for a signal. Fig.8(a) shows how the segmentation looks like across all four automata states from the pivot approach (each automata state is represented by a colored box, transitions occur based on selected criteria [3]).

B. Composites

The next layer was designed to extract action or motion compositions (MCs) by analyzing order-pair sequences of primitives. By studying the patterns in the ordered-pairs, action-level performance can be understood from the data. In total, seven basic MCs were derived: adjustment, increase, decrease, constant, contact, positive contact, negative contact, and unstable motions. Positive gradients regardless of magnitude were paired as a single group 'Positive' gradients while the same was done with 'Negative' gradients.

Table I summarizes motion compositions classification based on primitive ordered-pairs. The table contains subtables representing five primitive groupings. The first primitive is in bold followed by a listing of possible primitives with corresponding motion compositions and labels to classify an MC. An example of the MC layer can be visualized in Fig. 8(a).

C. Low-Level Behaviors

The taxonomy's third layer considers MC ordered pairs along with contextual information such as signal duration and amplitude values to yield classifications. Eight LLB classifications were derived and labeled as: push, 'PS', pull, 'PL', contact, 'CT', fixed, 'FX', alignment, 'ALIGN', shift, 'SH', and noise, 'N'. The LLB formulation criteria is similar

pimp	$m \geq 70$
bpos	$46 \leq m < 70$
mpos	$23 \leq m < 46$
spos	$1.0 \leq m\ 23$
const	$-1.0 < m < 1.0$
sneg	$-23 < m \leq -1.0$
mneg	$-46 < m \leq -23$
bneg	$-70 < m \leq -46$
nimp	$m \leq -70$

Fig. 4. Gradient values classification for the Primitives layers.

to those at the MC level. That is, for a pair of increase MCs labels, or decrease MCs labels, or constant MCs labels or adjust MCs labels; pull, push, fixed, or adjust LLBs are assigned respectively. As for contacts, if there is a positive contact followed by a negative one, or vice-versa, a contact LLB is assigned. One major difference between the MC level and the LLB level is introduction a shifting behavior 'SH'. Shifts and alignments are similar but differ in that, whenever there are two contiguous adjustment compositions, if the second composite's amplitude is larger than the first, label it as 'SH' LLB, if smaller label it 'ALGN' LLB (see[1] for more details).

D. Refinement

Refinement stages are critical to the system acting as filters for both the system's MC and LLB layers. Each refinement stage cleans up MCs and LLBs generated during the initial algorithm pass. Each refinement stage analyzes each layer on the basis of three contexts: a composite's duration, a composite's amplitude value, and repeated composites:
- Time Context Filter: examines two contiguous composites (except for contacts). If one composite is much larger than another one, the smaller composite is considered trivial and is absorbed into the larger composite (see [1] for further details).
- Amplitude Context: the amplitude context compares two types of values (average and amplitude values) for contiguous composites and considers whether the composites can be merged. Similar average values indicate that both composites belong to the same region. Similar amplitude values increase the likelihood of representing the same event. The criteria for MCs and LLBs is defined below.

For the MCs: (i) 'i-d' pairs in either ordered are merged into adjustments, (ii) 'i-k' or 'd-k' pairs in either order are merged as increases or decreases respectively,

For the LLBs: (i) 'PS-PL' pairs in either order are merged into 'ALIGN,' (ii) 'SH-ALIGN' in either order where the second composite has a smaller amplitude, are merged as 'ALIGN,' and (iii) 'ALIGN-PS||PL' or 'SH-PS||PL' or in either order are merged as an alignment'ALIGN.'
- Repeated Context: If two contiguous composites have repeated labels, they are merged as one, and their data structures updated accordingly.

The refinement cycle thus finds patterns across composites that are not evident when primitive or composite ordered-pairs are initially analyzed. Furthermore, as the refinement cycles are run at the different abstraction levels, refinements at higher abstraction levels filter hidden patterns at lower

HLB	LLB	Force Axis
Rotation	FX	Fx
	FX	Fz
	FX	My
Snap	CT	Fx, MY
Mating	AL || FX	Fy-Mx, Mz
	AL || FX	Fx-Mz

Fig. 5. Combinations of necessary LLBs to ascertain the presence of HLBs.

abstraction levels. The refinement filter was run three times in each layer to merge most disjoint composites. Fig. 8(a) already shows post-refinement results.

E. High-Level Behaviors

The fifth layer contextualizes state reasoning by asking: "What LLBs **principally** describe the automata Rotation, Snap, and Mating states?" The HLB characterizes what LLB configuration, across all six force axes, define a given HLB. Currently we only focus on defining the HLBs that yield an assembly successful (we will later focus on failure cases). The selection of key LLBs is connected with the controllers and reference parameter selection in the Pivot Approach as well as the local task coordinate frame selection for the task (see [2] for a more detailed description). In [1], an outcome assessment method was implemented based on the presence of successful LLBs and correspondingly HLBs. In [2] a bayesian filter was implemented to yield probabilistic beliefs about the LLBs and HLBs and it was accompanied by an outcome assessment scheme to determine to infer if a task was successful based on the beliefs. The table of Fig. 5 summarizes what LLBs characterize the above-stated desired HLBs:

V. Gradient Calibration

In the context of the PA and the RCBHT system, a gradient calibration routine was devised to acquire a set of gradient thresholds for the primitive layer that will be effective for a given robot-snap-part pair. That is, obtaining gradients that will optimally classify the FT signals into their appropriate labels the majority or totality of the time. As mentioned in Sec. IV-A, the most important labels in our gradient classification scheme are the ones used to classify *contacts* (which yield large changes in FT data) and *constant* actions or *fixed* behaviors (where there is little to no change in FT data). Both of these labels have a positive and negative version. A total of eight labels are used to classify the gradient space as described in Sec. IV-A: $pimp, bpos, mpos, spos, sneg, mneg, bneg, nimp$. In our previous works, these values were obtained by trial and error. But in this work we devised a scheme to calibrate these values such that they will be effective in appropriately classifying FT signals as long as the same robot is working with the same part. This statement assumes that the PA is utilized. This is important as the PA constraints the snap assembly motion in the same way across trial thus enabling for similar patterns, in form and in magnitude, of FT signals to be generated across trials.

A. Gradient Thresholds Determination

The first step as part of gradient calibration consists in determining the values for the contact gradients (labeled 'pimp' or 'nimp') and the constant gradients (between 'spos' and 'sneg'). Once these thresholds have been computed the space between 'spos' and 'bpos' and their negative counterpart can be divided into equally spaced segments, as in the table of Fig. 4. In our work, we evaluated values of gradients within specific automata states (further detailed in the next section). A gradient was considered to exist within an automata state if the beginning of the primitive (or linear segment) started after the automata state began and finished before the automata state terminated.

1) Contacts: For contact LLBs, as per our PA key LLB selection criteria, typically occur in the Snap state of the Fx and My axes, and they only occur a small number of times. Statistical measures like the mean, median, or mode are not useful to extract an effective contact threshold. Instead, the absolute value of the maximum gradient is used. Once the maximum gradient is found, the 'pimp' label in the primitive's layer is scaled by a constant k, such that: $pimp = k*(max(abs(grad_m)_{state_axis})), m \in M$, where M is the set of gradients in a given automata state in a given axis. Since this calibrated contact value will be used across

TABLE I
Motion Compositions according to Primitive Pairs

Combination	Category	Label	Combination	Category	Label	Combination	Category	Label
Positive			**Pimp**			**Constant**		
Negative	adjustment	a	Positive	pos contact	pc	Positive	increase	i
Positive	increase	i	Negative	pos contact	pc	Negative	decrease	d
Constant	increase	i	Constant	pos contact	pc	Constant	constant	k
Pimp	pos contact	pc	Pimp	unstable	u	Pimp	pos contact	pc
Nimp	neg contact	nc	Nimp	contact	c	Nimp	neg contact	nc
Negative			**Nimp**					
Positive	adjustment	a	Positive	neg contact	nc			
Negative	decrease	d	Negative:	neg contact	nc			
Constant	decrease	d	Constant	neg contact	nc			
Pimp	pos contact	pc	Pimp	contact	c			
Nimp	neg contact	nc	Nimp	unstable	u			

Axes	Fx	Fy	Mx	My	Mz
`pimp' axes	*Fx*	*Fz*	*My*	*My*	*My/10*

Fig. 6. Summary of 'pimp' threshold assignments across force axes as part of gradient contextualization.

trials, the value is scaled down to increase the likelihood of capturing contacts with similar values across trials. During the training phase an initial scaling factor of 0.90 was used but was later changed to 0.85 and will be described in the experiments.

2) Constants: For fixed LLBs, the mean, median, and mode values were compared to find which measure would yield the most effective threshold for constant signals. Experimental results (see Sec. V-B) identified the mean to be the best measure. Hence, the 'spos' label in the primitive layer was set to the absolute value of the gradient mean as: $spos = abs(mean(grad_m)_{state_axis})$.

3) Contextualization: Another important consideration lied in whether we should compute such thresholds for each separate axes and in some cases for each automata stare or not. As per our previous findings, a successful assembly can be characterized by select LLBs which capture that main components of the task. Similarly, it is the gradients found in the same automata states as the key LLBs that dominate the classifications. By contextualizing the gradient classification in this selective way, a more effective calibration can be attained. We begin with the contextualization of the contact thresholds and then proceed with the constant thresholds.

The Contact HLB is classified by two 'CT' LLBs in the Snap automata state along the Fx and My axes. For this threshold, the 'pimp' value of the Fz axis was used to class the Fz and Fy axes and the 'pimp' value of the Fx axis was used for itself. The reason we separated the classification in this way was because the average value of gradients for the Fy and Fz axes was approximately 31 while that of Fx was 81. The Fx axis in world coordinates is the axis in which an initial vertical contact takes place between the camera mold parts upon snapping and represents the hardest contact. For the moment axes, the 'pimp' value of My was used for itself and for Mx but not for Mz we used one-tenth of the value of My as the mean of the latter is an order-of-magnitude smaller. A summary of gradient contextualization can be seen in Table 6. Once all gradient thresholds could be derived for a given trial, those values would be used to test whether or not the RCBHT system would classify successful assemblies (as observed by appropriate snapping and mating) across a number of training assemblies. This part of the experimentation further allowed us to make observations that were included in our calibration approach. The training experiments are described in the following section.

B. Experiments

In our experiments, five training and seven test assemblies were run. Each of the training and test trials were successful as supervised by an external user. In the training session, gradient thresholds were computed for each of the five trials. Then, each trial was assessed by each of the five different gradient sets of thresholds as part of the RCBHT system to determine whether the task was successful or not. In effect, 25 runs were attempted to assess the trial outcomes. As described in [1] and [2], the system considers the task to be successful if all key LLBs are present in the Rotation, Snap, and Mating automata states. In this way, if the RCBHT declares the outcome to be successful we know that the gradient calibration was effective for that trial.

The results were organized in a table as shown in the table of Fig. 7 in order of increasing magnitude for the 'pimp' threshold of the trials. From these experiments we can immediately note that it is those trials that yielded the contact thresholds of lowest magnitudes that yielded a successful outcome for other trials. This is so since the contact labels in trials with higher thresholds remain being contact labels in those trials. However, when the contact threshold of a trial is greater than a threshold in another trial, then the contacts there cannot be discerned and thus the RCBHT will not encounter contacts in the trial which are necessary in order for the Snap HLB to exist. The trend is that as the 'pimp' threshold increases per trial, the likelihood of assessing an outcome as success diminishes.

There are a few exceptions to this trend and they are due to a number of different issues. In the table of Fig. 7, note that entries can be described by the *trial* number of the row and the *trial* number of the column. For entry (trial 1, trial 3), a failure might have been expected, however as noted in Sec. V-A.1, the 'pimp' value was scaled by a factor of 0.90 originally and later by a factor of 0.85. By lowering the scaling factor by 0.05 points, it allows some trials with higher contact thresholds to still identify contacts in trials with lower contact thresholds (in fact 2 out of 5 trials behaved this way).

For entry (trial 4, trial 2), another failure might have been expected. In this case, a contact LLB appeared in what may be considered the transitional period between states. That is, the 'CT' LLB started within the Snap state but ended within the MAT state. As described in Sec. V-A, we did not consider these kinds of transitional behaviors. This detected CT state for trial 2, had a high enough value that trials 4 and 5 detected it (along with other key LLBs) and rendered the outcome as successful. Dealing with transitional information is an important aspect that needs to be addressed. It was first

Training Trials	Max Grad	Three Clean Up Cycles					Succes Rate
		3	1	2	4	5	
3	60.32						100%
1	67.95						100%
2	92.48						40%
4	101.62						60%
5	105.24						40%

Fig. 7. Five training trials are listed in order of increasing contact threshold value. The green color indicates the RCBHT assessed the task as successful. The orange color indicates the RCBHT assessed the task as failure.

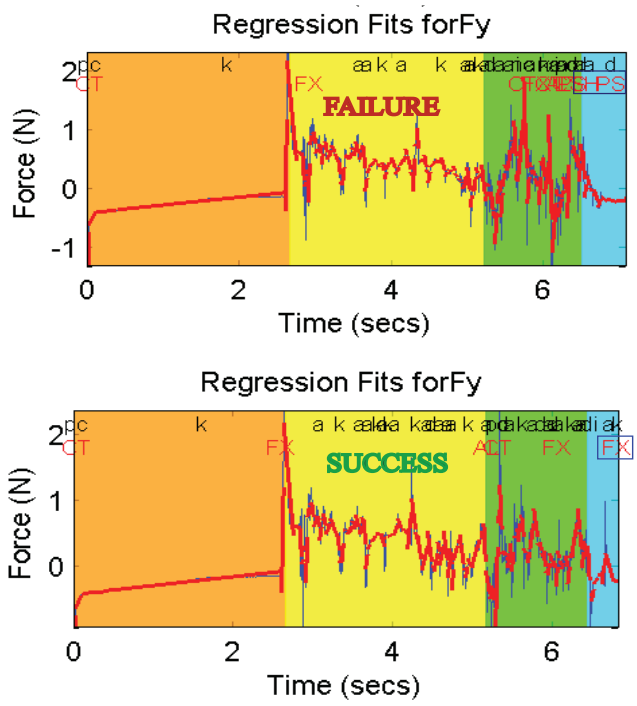

Fig. 8. The top figure results used 3 clean up cycles as part of the filter at the motion composition level. The three cycles may filter important compositions at the second level of the system. The bottom figure shows how when the clean up cycles are reduced to 2, the constant action (and correspondingly the 'FX' LLB) appears.

identified in [2] and will be addressed in the near future.

For entries (trial 2, trial 4) and (trial 5, trial 4), we have two failure assessments. The latter may have been expected but not the former. In fact, for entry (trial 5, trial 4) the scaling of the 'pimp' parameter would have allowed this trial to be assessed as successful, the problem lied elsewhere. No 'FX' behavior was identified during the Mating automata state in the Fy and Fz axes as can be seen in Fig. 8(a). The FT signal did indeed have a constant ('k') action at that stage, however, the latter had been absorbed by clean up process at the motion composition level. The latter used was absorbed into another action composition as part of a time context filtering after three clean up cycles. When, the filtering was reduced to two clean up cycles, the constant action appeared as well as the corresponding 'FX' LLB. This was not an isolated case as it also occurred in entry (trial 5, trial 4).

1) Minimum Gradient Selection: From the training experimentation, our calibration method was extended to consist of the following two aspects: (a) run at least five successful trial assemblies, compute their 'pimp' thresholds according to Sec. V-A, and choose the lowest magnitude threshold of the trial set; and (b) utilize two clean up filtering cycles at the motion composition level rather than three to avoid excessive filtering.

This gradient calibration methodology was run on seven test assemblies. The set of assemblies is also of interest because two of those trials presented false-positive results

Test Trials	Max Grad
1	56.45
2	57.68
3	91.89
4	80.00
5	73.69
6	74.83
7	76.56

Fig. 9. Seven test assemblies are listed with their contact threshold values.

when using the RCHBT in [2] (predicting a false outcome when in fact the assembly had been successful). The false-positive results had been a result of erroneous interpretation during the first two layers of the system. A summary of the contact gradient thresholds for the test assemblies is shown in the figure of Table 9. The first trial in the set yielded the smallest gradient and was used to generate the gradient classification thresholds that would be used for all test assemblies. For comparison the RCBHT was run once on the set of assemblies with two clean-up cycles and again another set with three cycles for comparison. The results are shown in the table of Fig. 10. The calibration method worked 100% of the time for the test assemblies set when using two clean up cycles suggesting to be an effective calibration method. It also worked 85.7% of the time when using three clean-up cycles. Fig. 11(a) shows how an uncalibrated system yielded a false-positive result in that a successful assembly did not posses a 'FX' LLB in the Rotation state of the My axes. In Fig. 11(b) on the other hand, the calibrated method was able to disambiguate the FT signal to show the presence of 'FX' LLBs. For further exploration, we tested the gradient thresholds from trial 1 in the training set with the two false-positive trials in the training set. In this case, when using two clean-up cycles both false-positive trials were correctly assessed as successful. When using three clean up cycles, it could only correctly interpret one of the two trials.

VI. Discussion

An effective gradient calibration method was implemented for the RCBHT system. Statistical measures were used to derive contact and constant gradient thresholds in contextually sensitive ways. During our experimentation we discovered that by running a number of training trials and selecting the gradient thresholds that belong to the trial whose contact threshold is the lowest, then the gradient calibration would be more effective. The experimentation also revealed a hidden

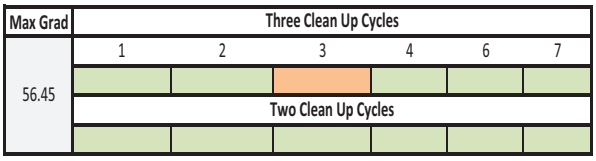

Fig. 10. Our calibration method with 2 clean up cycles successfully interpreted all seven test assemblies. Including two assemblies that had yielded false-positive results when using un-calibrated gradient thresholds.

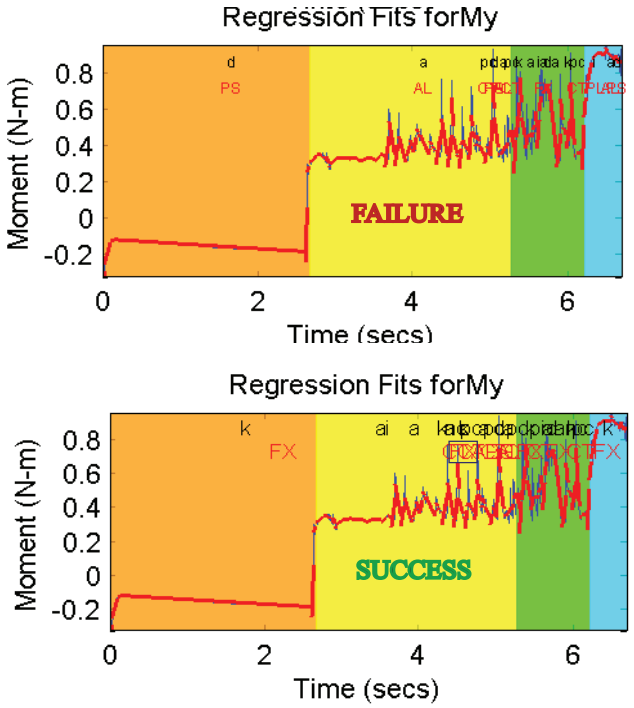

Fig. 11. The top figure shows how the RCBHT assigned a false-positive result by assigning a successive trial as failure. The bottom figures shows the difference in results after using the calibration method with two clean up cycles.

fact prior to this work. That is, that when using three rounds of filtering at the motion composition level, some key motion composition actions are filtered away by the system. Two clean up cycles proved to be a very effective number for correctly assessing snap assembly outcomes.

One of the keys in effectively calibrating gradients is understanding the assembly stages properly. The PA is useful in dividing the snap assembly into automata states that are consistent across trials. Each of those automata states have similar signal-patterns across trials for the six different force-torque axis. Furthermore, by having the RCBHT system abstract relative change across increasing layers of intuition, we can identify actions and behaviors that consistently appear across trials when the task is successful. Herein lies the significance of our framework; moreover, in this work's context, the design of the gradient calibration method exploits such interpretations in order to identify what to look for in the FT signals and then implement statistical measures to identify relevant thresholds. Our calibration method can be applied to new robots or different cantilever snaps if the same strategy and system are used since the strategy's docking action generates an pattern that is independent of robot or cantilever-part. This is crucial to increasing the system's viability to automate assemblies.

It is also worth mentioning, that the calibration method assumed a constant value for the 'determination coefficient' described in Sec. IV-A. The calibration should not be affected by different correlation values as long as they are constant throughout all trials. Another aspect worth noting is that the threshold values produced by the calibration method are tied to the FT transition values employed in our controller as part of the PA. If the FT transition conditions were to change, the RCBHT system would have to be recalibrated. Currently the FT transition values of the PA are also experimentally deduced. FT transition conditions in the PA can be found in [2], [1], [3].

One of the limitations of this work is the limited number of trial and test assemblies that were performed. By executing a higher number of trials our results will be more statistically significant. Similarly, all testing thus far has been performed solely on simulation. We intend to address both of these issues in the near future.

VII. Conclusion

In this work a gradient calibration method was presented as part of the Relative-Change-Based-Hierarchical Taxonomy (RCBHT) cantilever-snap verification system and the Pivot Approach control strategy for the automation of cantilever-snaps. As part of a relative-change based force signal interpretation scheme, an effective gradient calibration process is needed to increase the RCBHT's system robustness. Prior to this work, all gradient classification schemes were derived on an intuitive trial and error basis. Statistical measures were used to derive contact and constant gradient thresholds in contextually sensitive ways. The method requires training assemblies to identify a minimum contact gradient which serves as a marker for all other gradient thresholds. Experimental procedures verified that our calibration method was effective. Assemblies with supervised successful outcomes were used in experimentation. The RCBHT assessed out assemblies as successful using the calibration method. Even two snaps that where classified falsely as unsuccessful when using a previously non-calibrated version of the RCBHT.

References

[1] J. Rojas, K. Harada, H. Onda, N. Yamanobe, E. Yoshida, K. Nagata, and Y. Kawai, "A relative-change-based hierarchical taxonomy for cantilever-snap assembly verification (in print)," in *IEEE Intl Conf. on Robots and Systems*, 2012.

[2] ——, "Gradient calibration for the rcbht cantilever snap verification system." in *IEEE-RAS Intl. Conf. on Humanoid Robots*, 2012.

[3] ——, "Cantilever snap assembly automation using a constraint-based pivot approach," in *IEEE Intl. Conf. on Mechatr. & Automation*, 2012.

[4] F. Kanehiro, H. Hirukawana, and S. Kajita, "Openhrp: Open architecture humanoid robotics platform," *Intl. J. of Robotics Res.*, vol. 23-2, pp. 155–165, 2004.

[5] H. Asada, "Teaching and learning of compliance using neural nets: Representation and generation of non-linear compliance," in *IEEE Int'l Conf. on Robotics and Automation*, 1990.

[6] B. McCarragher, "Force sensing from human demonstration using a hybrid dynamical model and qualitative reasoning," in *IEEE Intl. Conf. on Robotics and Automation*, 1994, pp. 557–563.

[7] W. Meeussen, J. Rutgeerts, K. Gadeyne, H. Bruyninckx, and J. D. Schutter, "Contact-state segmentation using particle filters for programming by human demonstration in compliant-motion tasks," *IEEE Trans. on Robotics*, vol. 23.2, pp. 218–231, 2007.

[8] K. Gadeyne, T. Lefebvre, and H. Bruyninckx, "Bayesian hybrid model state estimation applied to simultaneous contact formation recognition and geometrical parameter estimation," *Int. J. Robotics Research*, vol. 24(8), p. 615630, 2005.