

XXIX IFSO WORLD CONGRESS 2026 · TORONTO

Rapid Proficiency without a sustained learning curve

Console-time dynamics in robotic RYGB on the Hugo™ RAS
platform

I. Mesquita, F. Rosas, JP Santos, T. Correia, M. Marcos, J. Santos, P. Soares

Bariatric, Metabolic & Upper-GI Surgery · Unidade Local de Saúde de Santo António, Porto, Portugal

Disclosures: proctor for Medtronic (Hugo™ RAS). No funding was received for this study.



Isabel Mesquita, MD FACS, FEBS MIS RAS



Disclosures

No conflicts of interest to declare.

Funding	No industry funding, consultancy or speaker fees related to this work.
Ethics	Retrospective analysis approved by the institutional ethics committee.
Data	Console time machine-recorded by the platform; no surgeon-reported timings.

ULS Santo António

A tertiary metabolic bariatric unit in Porto, operating the Hugo™ RAS platform since March 2023.

210 robotic metabolic bariatric procedures to June 2026.

The “cases-to-plateau” paradigm

How robotic learning curves are conventionally reported

01 A fixed number of cases

Reported as a single threshold — the case at which operative time plateaus. Published values vary widely between series.

02 Derived from a different population

Most curves come from surgeons early in their robotic — and often bariatric — experience, on one platform generation.

03 Applied to everyone

The same threshold then governs credentialling, proctoring and programme planning, regardless of prior expertise.

OUR QUESTION

Does a sustained learning curve exist when experienced metabolic bariatric surgeons adopt a new robotic platform?

Hugo™ RAS — modular, open-console architecture, introduced 2022

Methods

Design

Single-centre retrospective cohort
ULS Santo António, Porto
March 2023 – June 2026

Population

210 metabolic bariatric procedures
on the Hugo™ RAS platform
by a four-surgeon team

Analytical cohort

91 primary RYGB with
machine-recorded console time
(three surgeons)

Analysis


Console time as primary measure
CUSUM · change-point regression
Bootstrap 95% CI

Why primary RYGB, and why three surgeons

- RYGB is the most expressive primary procedure in our programme — the only one with the per-surgeon volume a curve analysis needs.
- RYGB activity is concentrated in three surgeons, with machine-recorded console time for all their cases — n = 22, 40 and 29, together the 91 analysed.
- The fourth surgeon contributes to the programme denominator only

From programme to analytical cohort

All robotic metabolic bariatric activity, March 2023 – June 2026



Console time

Median (IQR), machine-recorded

85 min

OVERALL MEDIAN
IQR 71 – 101 min

Surgeon IM

81 min

n = 22 (47 *)

Surgeon MM

89 min

n = 40 (70)

Surgeon PS

90 min


n = 29 (71)

9 minutes separate the three surgeon's median console times — across an entire adoption period.

Times comparable to established laparoscopic and mature robotic RYGB practice were reached within the earliest cases of each series.

Console-time trend per surgeon

Individual cases in chronological order, with a centred 5-case moving median



Chronological order by verified operative date. Moving median centred on each case and defined from case 3 onward.

A short early phase, then a stable plateau. The first five cases run longest for every surgeon — median 85, 134 and 121 minutes for IM, MM and PS — falling to 61, 94 and 79 minutes over the next five. The descent is short and early; the scatter beyond it is dispersed throughout the series rather than concentrated at its start.

How brief is the early phase?

Console time, cases 1–5 versus subsequent

Cases 1–5 vs subsequent

121 → 82 min

$p < 0.001$

Same, excluding each surgeon's longest case

110 → 82 min

$p = 0.005$

From case 6 onwards, trend vs case number

$\rho +0.01$

$p = 0.92$

Pooled across the three surgeons: 15 early cases versus 76 subsequent. Mann–Whitney; Spearman for the trend.

Within the first cases. Beyond it the series are flat.

BUT “FIVE” IS NOT A THRESHOLD

It was not estimated.

The figure comes from truncating the series case by case until the trend disappears — not from a model.

The model could not find it.

Change-point regression returned a bootstrap 95% CI spanning the whole series.

But the finding is robust to where you cut.

Any cut between cases 3 and 8 gives $p \leq 0.006$, and the rest of the series stays flat ($p -0.12$ to $+0.08$, all NS).

CUSUM and change-point analysis

An early transition, consistent across surgeons — but not localisable




CUSUM

A short ascending phase, then sustained descent in all three series. The peak falls at case 1, 8 and 4 for IM, MM and PS.


CHANGE-POINT REGRESSION

Point estimates agree — case 3, 4 and 3 — but the bootstrap 95% CI spans 64–82% of each series.

The transition is real and early. Its position cannot be pinned to a single case number — so no threshold should be quoted from it.



Change-point estimate with bootstrap 95% CI



CUSUM = cumulative deviation from each surgeon's own mean console time. Shaded band = bootstrap 95% CI (3000 resamples); vertical mark = point estimate.

Safety was not traded for speed

30-day outcomes in the 91 analysed primary RYGB

1.1%

CONVERSIONS TO
LAPAROSCOPY / OPEN

one in the analysed cohort

7.7%

MAJOR COMPLICATIONS

7 of 91 · Clavien–Dindo $\geq$ IIIa

0

30-DAY
MORTALITY

no deaths in the series

2 days

MEDIAN
LENGTH OF STAY

IQR 1 – 3 days

No early-series clustering

Complications fell at cases 7, 13, 20 and 35 for MM and 3, 10 and 13 for PS; IM had none. Pooled, 1 of the first 15 cases (6.7%) versus 6 of the subsequent 76 (7.9%), $p = 1.00$; extending to the first ten cases per surgeon, 3 of 30 versus 4 of 61, $p = 0.68$.

Denominators are the 91 analysed cases, not the 210-procedure programme. Three primary bypasses converted to laparoscopy or open surgery fall outside this cohort: one by the fourth surgeon, two excluded by the console-time filter. Fisher exact test.

Why is the curve so short?

Four readings of a brief early phase in experienced laparoscopic surgeons

Expertise transfers

The cognitive and procedural core of RYGB — the anatomy, the sequence, the judgement — is already mastered. What is left to learn is the interface, and it is learned quickly.

Platform design may support rapid adaptation

Hugo™ RAS combines a modular arm architecture with an open console. These features may ease adaptation for experienced laparoscopic surgeons. Platform-specific set-up — docking, arm positioning and port geometry — also stabilised early.

Beyond the first few cases, variability is not sequence

The long cases from then on were dispersed throughout each series rather than concentrated at its start ($p = 0.01$, $p = 0.92$).

Published thresholds measure something else

Case-to-plateau figures of 50 to 100 come from series that pool surgeons into a single curve, that include surgeons early in their bariatric careers, and that measure total operative time — which carries the team's learning, not only the surgeon's. Ours is per-surgeon, on machine-recorded console time.

Everyone converges — but not everyone starts

Early phase and stable level, by surgeon

Surgeon MM

Opened the programme, March 2023

134 → 81_{min}

cases 1–5 → stable level

Early-phase excess **+66% p = 0.002**

Stable level reached **case #7**

Surgeon PS

Opened the programme, March 2023

121 → 86_{min}

cases 1–5 → stable level

Early-phase excess **+40% p = 0.009**

Stable level reached **case #6**

Surgeon IM

Joined 12 months later

85 → 79_{min}

cases 1–5 → stable level

Early-phase excess **+7% p = 0.28**

Stable level reached **case #2**

79, 81 and 86 minutes.

Three surgeons, three different starting points, a seven-minute spread at the end. The two who opened the programme each worked through a measurable early phase; the one who joined eleven months later, into a settled technique and team, reached the stable level by her second case. Programme maturity appears to shorten the curve for those who follow — on a single observation.

Early phase = cases 1–5; stable level = median of cases 6 onwards. Mann–Whitney per surgeon. “Stable level reached” = first case at or below that median.

Limitations

01 Console time is a proxy, and was recorded twice

It measures efficiency, not the quality of the anastomosis or of patient selection. Two time fields disagreed in 8 cases; the raw timestamp was used, as the derived field exceeded total operative time.

02 The cohort is selected

91 cases, one high-volume unit, retrospective. The fourth surgeon's 22 cases carry no console time. Three converted primary bypasses fall outside the cohort — and conversion is exactly the event a learning curve should capture.

03 Complexity was not adjusted for

Case difficulty is invoked as an explanation but not modelled, although BMI and total operative time were available for all 91. Non-console time was described but recorded in only 69.

04 The programme-maturity finding is a single observation

One surgeon joined an established unit — and she also brought prior robotic upper-GI experience. Programme effect and personal experience cannot be separated here.

Ongoing: date verification against the clinical record, multivariable modelling of case complexity and of docking / non-console time, and prospective validation across centres.

Conclusion

Experienced surgeons reached their stable console time within the first few cases on the Hugo™ RAS platform.

- The early learning effect was real but short, with no reproducible case-number threshold.
- Three surgeons started at different points but converged on similar console times.
- Safety was maintained from the beginning.
- Rapid adaptation may reflect transferred surgical expertise, platform design and programme maturity.
- These findings support structured platform adoption rather than one fixed case number for everyone.

XXIX IFSO WORLD CONGRESS 2026 · TORONTO

Rapid Proficiency without a sustained learning curve

Console-time dynamics in robotic RYGB on the Hugo™ RAS
platform

I. Mesquita, JP Santos, F. Rosas, T. Correia, M. Marcos, J. Santos, P. Soares

Bariatric, Metabolic & Upper-GI Surgery · Unidade Local de Saúde de Santo António, Porto, Portugal

Disclosures: proctor for Medtronic (Hugo™ RAS). No funding was received for this study.

Thank you

Questions welcome

Isabel Mesquita, MD, FACS, FEBS MIS RAS



XXX IFSO WORLD CONGRESS

"Advancing Multimodal Obesity Care"

24-27 August 2027

ifso2027.org



IFSO
PARIS 2027

See you in

PARIS



Console time by surgeon — detail

	n	Median (IQR)	Range	First 10	Subsequent	p
Surgeon IM	22	81 (67–98)	52–165	91	83	0.51
Surgeon MM	40	89 (74–109)	42–244	100	87	0.11
Surgeon PS	29	90 (72–101)	44–341	96	86	0.13

Console time in minutes. “First 10” and “Subsequent” are medians; p from Mann–Whitney U.

The first-10 versus subsequent difference is 8–14 minutes and reaches significance in none of the three surgeons individually — the effect is concentrated in the first five cases, where the earlier slide locates it.

Where are the long cases?

Nine cases exceeded the 90th percentile (≥ 135 min). Their position within each surgeon's series:



3 of 9

of the longest cases occurred in the second half of a surgeon's series — and the two extremes were case 2 and case 3.

BMI alone did not predict console time ($p = 0.08$, $p = 0.43$).

And note what is not in the cohort: eight primary bypasses with an added step — remnant gastrectomy, cholecystectomy — were excluded as combined procedures. The technically heaviest cases were removed.

Anticipated questions

“So the learning curve is five cases.”

That is a fair reading, and we do not resist it — cases 1–5 took 106 min against 84 thereafter. Our point is that five is a truncation, not an estimate: the change-point model could not locate it, and a plateau that early may be unidentifiable in principle.

“Why exclude the fourth surgeon?”

Console time was never recorded for his cases. Including him would have meant substituting total operative time, which conflates docking and non-console phases.

“Isn’t 91 cases too few to detect a change point?”

Power is a fair concern — but the individual series are the same size as those from which published thresholds are derived, where the CI is rarely reported. The difficulty here is the position of the change point, not only the sample size.

“What about outcomes beyond time?”

1.1% conversion, no mortality, median stay of 2 days, with no early-series clustering. Longer-term weight and metabolic outcomes are being analysed separately.