Supplemental Materials to *Control of speaking rate is achieved by switching between qualitatively distinct cognitive 'gaits': Evidence from simulation*

Joe Rodd[a b], Hans Rutger Bosker[a c], Mirjam Ernestus[b a], Phillip M. Alday[a], Antje S. Meyer[a c], and Louis ten Bosch[b a]

[a]*Max Planck Institute for Psycholinguistics, Nijmegen, the Netherlands*
[b]*Radboud University, Centre for Language Studies, Nijmegen, the Netherlands*
[c]*Radboud University, Donders Institute for Brain, Cognition and Behaviour, Nijmegen, the Netherlands*

Supplement A: Control model variant

Alongside the synchronous, asynchronous and linearly-constrained asynchronous model variants, we fitted a fourth model variant, the *control model variant*, that was most faithful to the original formulation in Dell, Burger, and Svec (1997). The performance of this model variant is, for unsurprising reasons, notably worse than the other variants, so we do not report it in the article text. For completeness, we report it here, and provide comparisons with the synchronous and asynchronous model variants.

The control model variant has a single duration parameter, dur_0, that controls the duration of all the phases of the activation pattern of both the first and second port. Each phase lasts dur_0 model "ticks" (an arbitrary time unit equivalent to 9ms). The frame output pattern of this model is depicted in the bottom cell of Figure 1, which complements Figure 3 in the article text. The control model variant inherits the overly simplistic assumption that the first and second syllables should be active for the same length of time.

Learning

The control model was fitted in the same way as the other models, as described in the article text in the Section headed "Optimization procedure". We used loess-fitting (Cleveland & Devlin, 1988) to identify the trend in the score for each objective function in each rate condition. These loess-fits are shown in the Figure 2, which complements Figure 6 in the article text. Relative to the synchronous and asynchronous model variants, the control model variant makes slower progress, and achieves notably poorer scores.

Convergence

To assess convergence, we calculated the normalised hyper-volume indicator (Zitzler, Brockhoff, & Thiele, 2007), and also the relegation count. These metrics are shown for the control, synchronous and asynchronous model variants in Figure 3, which complements Figure 7 in the article text.

Qualitative model performance

Figure 4 is a companion to Figure 9 in the article text. We show the distributions resulting from combining the duration distributions predicted by each member of the Pareto front of each run as solid violins. These are compared against the target distributions measured from the corpus (translucent violins with dashed edges).

The control model variant fits the data strikingly poorly, suffering from the same bi-modality problem as the synchronous model variant, but to a larger extent.

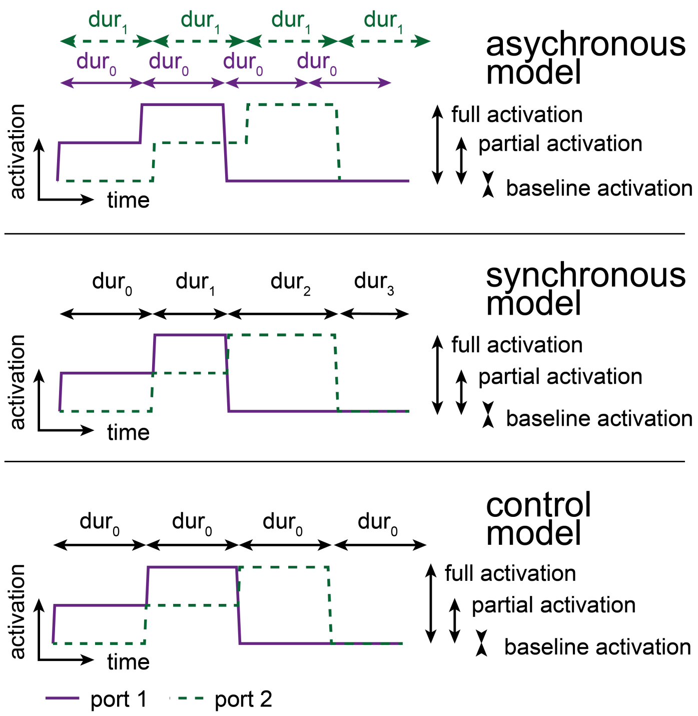

Figure 1: The activation patterns produced by the frame node for port 1 (purple, solid) and port 2 (green, dashed) in the asynchronous, synchronous and control models. The duration of each step in the activation patterns is controlled by various parameters, depending on the model variant (such as dur_0, see text for full details).

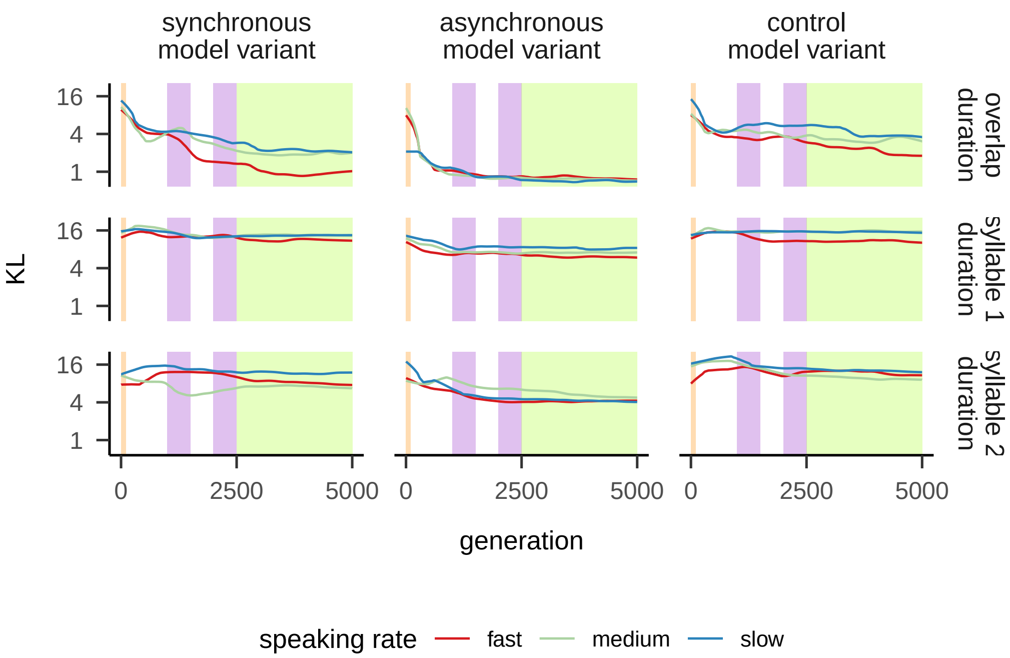

Figure 2: Loess-fits of the Kullback Leibler scores (y-axis, log-transformed scale, lower values indicate better performance) of the solutions in the Pareto front in each generation (x-axis), for the three rate conditions (line colours), the three objective functions (rows) and the synchronous, asynchronous and control model variants (columns). The shading indicates the optimisation phases of the model, orange is the phase where only the µ component of a subset of the parameters was adjusted by the optimiser, white indicates that the µ compoment of all parameters was adjusted by the optimiser, purple indicates that the σ component of all parameters was adjusted by the optimiser, and green indicates that both µ and σ components of all parameters were adjusted.

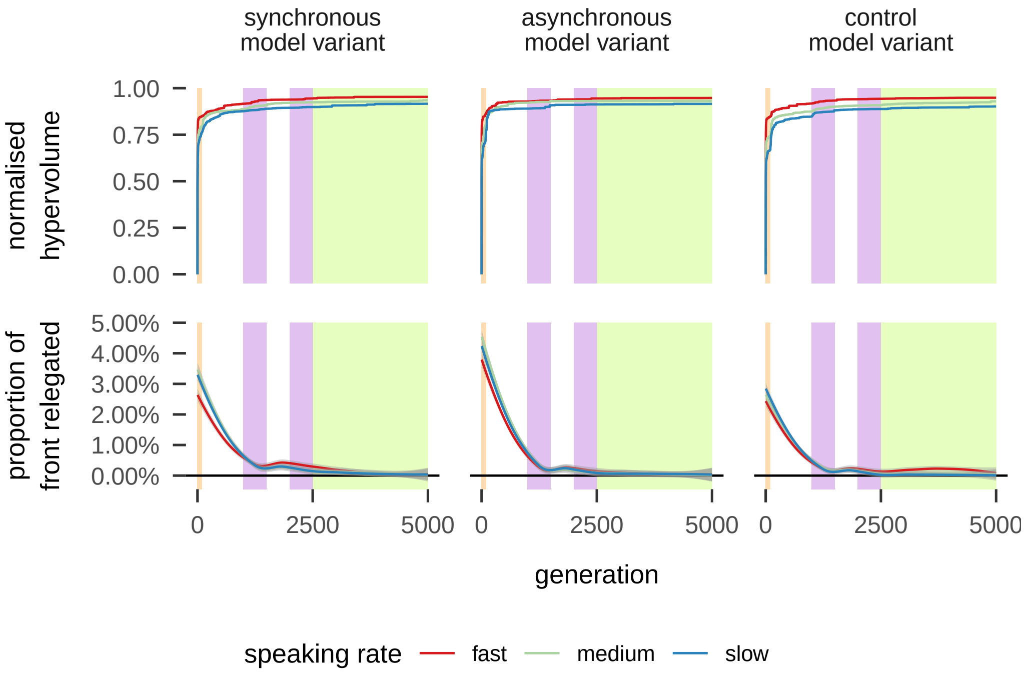

Figure 3: Upper panels: The normalised hypervolume indicator (y-axis) during the 5000 generations of the optimisation run (x-axis), for the three model variants (columns). Stabilisation of the normalised hypervolume indicator at a value close to 1.0 indicates successful convergence. The colour of the lines indicate the speech rate condition being optimised. Lower panels: the proportion of former front members relegated from the front in each generation. See the caption of Figure 2 for the meaning of the shading.

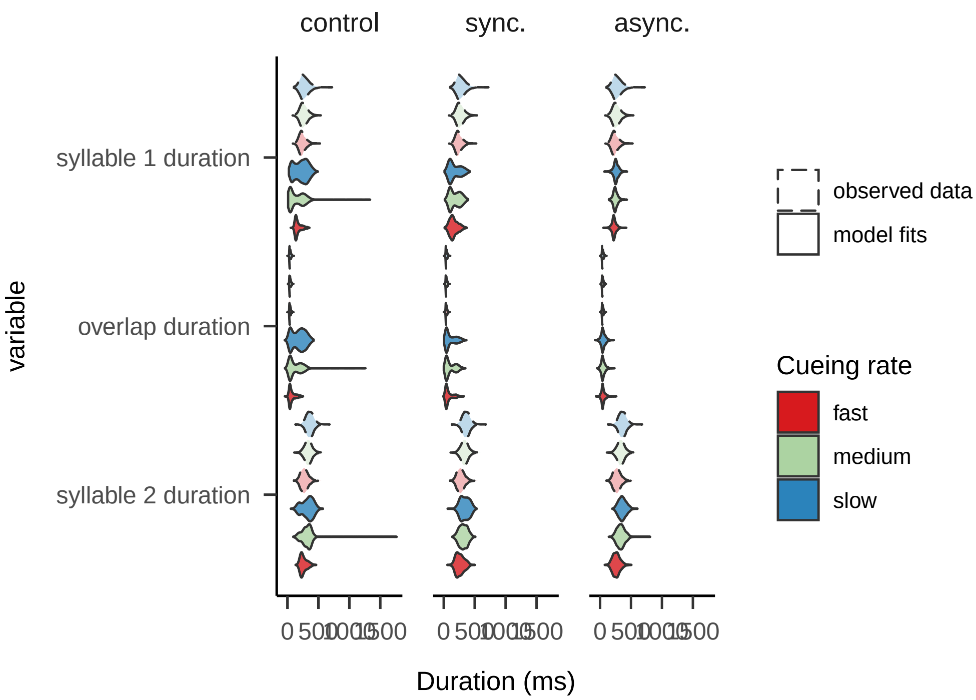

Figure 4: The duration (x-axis) distributions (filled violins) predicted by three models variants (facets) at the three rate conditions (colours) for each of the three target distributions (y-axis), compared against the observed distributions (translucent violins with dashed edges).

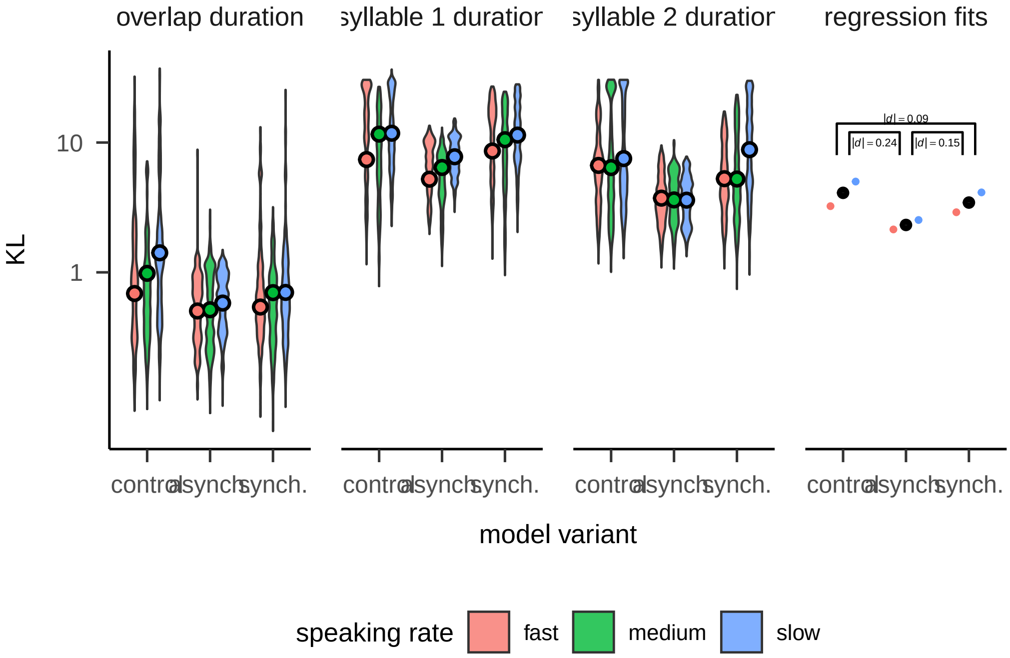

Figure 5: First three panels: the bootstrapped distributions (violins) of the KL scores (y-axis, smaller is better, log scale) achieved by the 0-ranked agents (the Pareto front) for each model variant (x-axis, *control*: control model variant, *async.*: asynchronous model variant, *sync.*: synchronous model variant). in each speaking rate condition (fill colours), in each objective (panels). The coloured dots indicate the model fits for the three-way interaction term in the regression model. Fourth panel: the fits of the model variant term from the regression model (main effect shown as black dots, fits of rate condition:model variant interaction in smaller coloured dots). 95% confidence intervals are omitted because they are too small to be visible. Significant differences in the main effect are indicated. The main effect of model variant is plain to see; the asynchronous model variant performs significantly better (achieves lower KL scores) than the synchronous and control model variants. In turn, the synchronous model variant outperforms the control model variant.

Statistical test of model performance

To assess the performance of the control model, the same statistical analysis was conducted as for the other comparisons. Figure 5 complements Figure 8 in the article text.

The regression table in Table 1 is comparable with the regression tables reported in Supplement C.

Table 1: Summary of regression model for log KL, comparing control, asynchronous and synchronous model variants

	Estimate (KL scale)	Estimate (log transformed)	Std. Error (log transformed)	t value	Pr(>|t|)
(Intercept)	0.25	-1.37	0.0008	-1670.37	< 0.001
Model variant: sync.	1.25	0.22	0.0010	215.15	< 0.001
Model variant: control	1.60	0.47	0.0010	451.51	< 0.001
Rate condition: fast	0.98	-0.02	0.0012	-13.78	< 0.001
Rate condition: slow	1.09	0.09	0.0012	71.99	< 0.001
Objective: syllable 1 duration	6.33	1.85	0.0012	1587.18	< 0.001
Objective: syllable 2 duration	4.15	1.42	0.0012	1225.21	< 0.001

Table 1: Summary of regression model for log KL, comparing control, asynchronous and synchronous model variants *(continued)*

	Estimate (KL scale)	Estimate (log transformed)	Std. Error (log transformed)	t value	Pr(>|t|)
Interaction: sync and fast	0.84	-0.17	0.0015	-116.67	< 0.001
Interaction: control and fast	0.78	-0.25	0.0015	-166.66	< 0.001
Interaction: sync and slow	0.92	-0.08	0.0015	-56.83	< 0.001
Interaction: control and slow	1.20	0.18	0.0015	118.56	< 0.001
Interaction: sync and syllable 1 duration	1.15	0.14	0.0015	94.92	< 0.001
Interaction: control and syllable 1 duration	0.96	-0.04	0.0015	-27.08	< 0.001
Interaction: sync and syllable 2 duration	1.05	0.05	0.0015	33.87	< 0.001
Interaction: control and syllable 2 duration	0.95	-0.05	0.0015	-35.75	< 0.001
Interaction: fast and syllable 1 duration	0.87	-0.14	0.0016	-83.66	< 0.001
Interaction: slow and syllable 1 duration	1.05	0.05	0.0017	30.69	< 0.001
Interaction: fast and syllable 2 duration	1.04	0.04	0.0016	24.00	< 0.001
Interaction: slow and syllable 2 duration	0.91	-0.09	0.0017	-52.64	< 0.001
Interaction: sync, fast, syllable 1	1.19	0.18	0.0021	85.06	< 0.001
Interaction: control, fast, syllable 1	1.07	0.07	0.0021	32.84	< 0.001
Interaction: sync, slow, syllable 1	1.00	0.00	0.0021	2.28	0.022
Interaction: control, slow, syllable 1	0.74	-0.31	0.0021	-143.46	< 0.001
Interaction: sync, fast, syllable 2	1.16	0.15	0.0021	73.27	< 0.001
Interaction: control, fast, syllable 2	1.29	0.25	0.0021	120.28	< 0.001
Interaction: sync, slow, syllable 2	1.60	0.47	0.0021	223.62	< 0.001
Interaction: control, slow, syllable 2	0.94	-0.06	0.0021	-27.81	< 0.001

Supplement B: Elicitation materials

Tables 2 and 3 list the words included in the production component of the study.

Filler words

Table 2: Filler words were included in the first, penultimate and last slots of each trial.

orthography	phonetic form	meaning	
gieter	ˈxi.tər	watering can	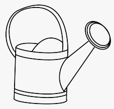
kabel	ˈkaː.bəl	cable	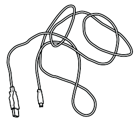
lasser	ˈlɑ.sər	welder	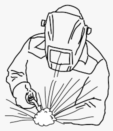
lichaam	ˈlɪx.aːm	body	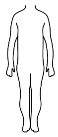
molen	ˈmoː.lən	windmill	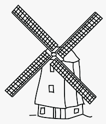
monnik	ˈmɔ.nɪk	monk	
spiegel	ˈspi.xəl	mirror	
tafel	ˈta.fəl	table	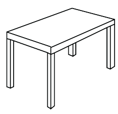
trommel	ˈtrɔ.məl	drum	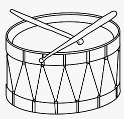

Table 2: Filler words were included in the first, penultimate and last slots of each trial. *(continued)*

orthography	phonetic form	meaning	
vinger	ˈza.ŋər	finger	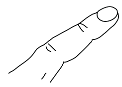
zanger	ˈza.ŋər	singer	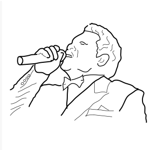

Target words

Table 3: Target words were included in the second to sixth slot of each trial.

orthography	phonetic form	meaning	
hagel	ˈhaː.xəl	hail	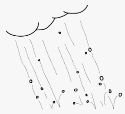
hamer	ˈhaː.mər	hammer	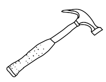
havik	ˈhaː.mər	hawk	
nagel	ˈnaː.xəl	fingernail	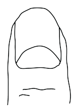
navel	ˈnaː.vəl	navel	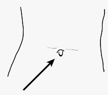
sinus	ˈsi.nʉs	sine wave	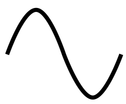
slager	ˈslaː.xər	butcher	

Table 3: Target words were included in the second to sixth slot of each trial. *(continued)*

orthography	phonetic form	meaning	
snavel	ˈsnaː.vəl	beak	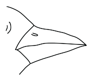
visum	ˈvi.sum	visa	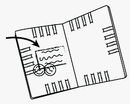
vlieger	ˈvli.xər	kite	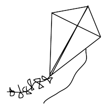
vriezer	ˈvri.zər	freezer	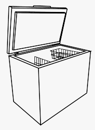
wafel	ˈwaːˌfəl	waffle	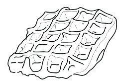
zoemer	ˈzu.mər	alarm	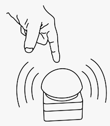

Distribution of target words in the dataset

Figure 6 shows the proportions of the three speaking rate sections of the corpus that each word represents. The fast speaking rate section shows more variation than the medium and slow speaking rate sections, but the overall range of proportions is relatively small, extending from 6.8% to 8.6% percent.

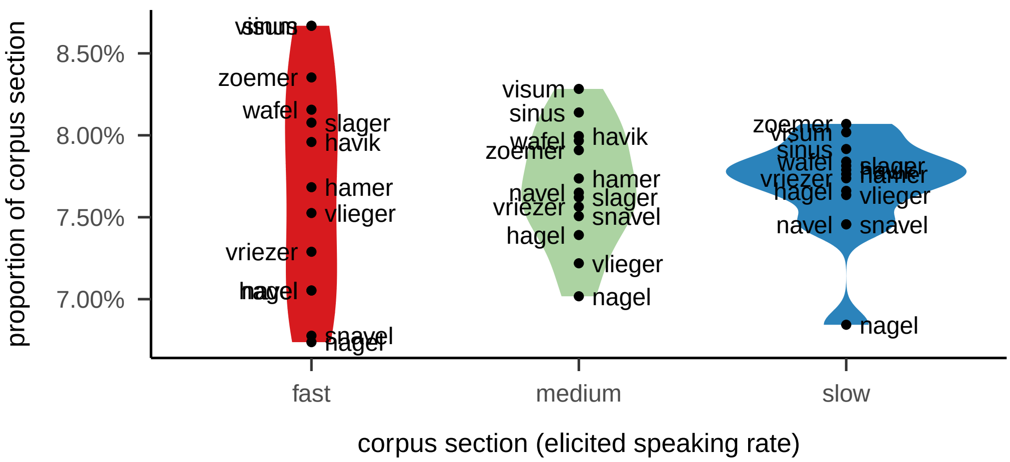

Figure 6: Distribution of target words in the dataset.

Supplement C: Tables for regression analysis

Modelling approach

In order to evaluate the performance of the different model variants, we need to identify and statistically test differences in the KL scores achieved by the Pareto front solutions of each of the model variants. Simultaneously, we need to disregard variation in the KL scores as a function of objective, since KL scores for the various objectives are not directly arithmetically comparable because of differences in the observed distributions, as previously discussed. The same holds for comparing models fitting different rate conditions, between which there are also differences in the variability of the observed distributions.

Instead of averaging scores across objectives, linear regression with categorical predictors for model variant, rate condition and objective can be used to isolate the effect on the KL score attributable to model variant, independent of rate condition and objective. This leads to a regression model with the following structure (R syntax):

```
log(KL) ~ model variant * rate condition * objective
```

This is a model predicting KL with categorical predictors for *model variant*, *rate condition* and *objective*, and all interactions between the levels of those categorical predictors.

The KL scores were bootstrap re-sampled to introduce variation required to perform regression modelling. The bootstrapped distributions of the KL scores are shown in the first three panels of Figure 8 in the article text. We took 2,000 samples with replacement of sets of syllable 1 duration, syllable 2 duration, and overlap duration values from the observed dataset. For each of these samples, we calculated the KLs between the re-sampled observed distributions and the model's predicted distributions. The resulting bootstrapped KLs were then log transformed and z-normalised. The log transformation was necessary to de-skew the KLs, which obey a log distribution.

Synchronous model variant vs. asynchronous model variant

This model compares the performance of the synchronous model variant vs. the asynchronous model variant. The latter is on the intercept (treatment coding). The response variable was the KL score of the relevant comparison between fitted distribution and observed distribution, so smaller values indicate better performance (less divergence between fitted and observed distribution). Table 4 summarises the model fit.

Table 4: Summary of regression model for log KL, comparing asynchronous and synchronous model variants

	Estimate (KL scale)	Estimate (log transformed)	Std. Error (log transformed)	t value	Pr(>|t|)
(Intercept)	0.27	-1.30	0.0007	-1802.03	< 0.001
Model variant: sync.	1.25	0.22	0.0009	247.12	< 0.001
Rate condition: fast	0.98	-0.02	0.0010	-15.83	< 0.001
Rate condition: slow	1.09	0.09	0.0011	82.69	< 0.001
Objective: syllable 1 duration	6.46	1.87	0.0010	1823.06	< 0.001
Objective: syllable 2 duration	4.22	1.44	0.0010	1407.31	< 0.001
Interaction: sync and fast	0.84	-0.17	0.0013	-134.01	< 0.001

Table 4: Summary of regression model for log KL, comparing asynchronous and synchronous model variants *(continued)*

	Estimate (KL scale)	Estimate (log transformed)	Std. Error (log transformed)	t value	Pr(>|t|)
Interaction: sync and slow	0.92	-0.08	0.0013	-65.28	< 0.001
Interaction: sync and syllable 1 duration	1.15	0.14	0.0013	109.03	< 0.001
Interaction: sync and syllable 2 duration	1.05	0.05	0.0013	38.91	< 0.001
Interaction: fast and syllable 1 duration	0.87	-0.14	0.0014	-96.10	< 0.001
Interaction: slow and syllable 1 duration	1.05	0.05	0.0015	35.25	< 0.001
Interaction: fast and syllable 2 duration	1.04	0.04	0.0014	27.56	< 0.001
Interaction: slow and syllable 2 duration	0.91	-0.09	0.0015	-60.46	< 0.001
Interaction: sync, fast, syllable 1	1.19	0.18	0.0018	97.70	< 0.001
Interaction: sync, slow, syllable 1	1.00	0.00	0.0018	2.62	0.009
Interaction: sync, fast, syllable 2	1.17	0.15	0.0018	84.16	< 0.001
Interaction: sync, slow, syllable 2	1.60	0.47	0.0018	256.85	< 0.001

Asynchronous model variant vs asynchronous model variant with linearity constraint

This model compares the performance of the asynchronous model variant with linearity constraint vs. the asynchronous model variant without linearity constraint. The latter is on the intercept (treatment coding). Table 5 summarises the model fit.

Table 5: Summary of regression model for log KL, comparing the asynchronous model variant and the asynchronous model variant with linearity constraint

	Estimate (KL scale)	Estimate (log transformed)	Std. Error (log transformed)	t value	Pr(>|t|)
(Intercept)	0.22	-1.51	0.0008	-1969.17	< 0.001
With linearity constraint	4.31	1.46	0.0014	1055.85	< 0.001
Rate condition: fast	0.98	-0.02	0.0011	-18.11	< 0.001
Rate condition: slow	1.11	0.11	0.0011	94.61	< 0.001
Objective: syllable 1 duration	9.62	2.26	0.0011	2085.92	< 0.001

Table 5: Summary of regression model for log KL, comparing the asynchronous model variant and the asynchronous model variant with linearity constraint *(continued)*

	Estimate (KL scale)	Estimate (log transformed)	Std. Error (log transformed)	t value	Pr(>|t|)
Objective: syllable 2 duration	5.74	1.75	0.0011	1610.22	< 0.001
Interaction: with constraint and fast	1.03	0.03	0.0019	13.66	< 0.001
Interaction: with constraint and slow	0.91	-0.10	0.0020	-48.26	< 0.001
Interaction: with constraint and syllable 1 duration	0.15	-1.92	0.0020	-979.74	< 0.001
Interaction: with constraint and syllable 2 duration	0.42	-0.87	0.0020	-443.45	< 0.001
Interaction: fast and syllable 1 duration	0.85	-0.17	0.0015	-109.95	< 0.001
Interaction: slow and syllable 1 duration	1.07	0.06	0.0016	40.34	< 0.001
Interaction: fast and syllable 2 duration	1.05	0.05	0.0015	31.54	< 0.001
Interaction: slow and syllable 2 duration	0.90	-0.11	0.0016	-69.18	< 0.001
Interaction: with constraint, fast, syllable 1	0.92	-0.08	0.0028	-29.71	< 0.001
Interaction: with constraint, slow, syllable 1	1.40	0.34	0.0028	121.01	< 0.001
Interaction: with constraint, fast, syllable 2	0.84	-0.18	0.0028	-63.68	< 0.001
Interaction: with constraint, slow, syllable 2	1.11	0.11	0.0028	37.69	< 0.001

Supplement D: Model parameters

The model has parameters that control many aspects of its behaviour and performance. These are illustrated in Figure 7, and listed with descriptions of their functions in Table 6. Some parameters are adopted directly from the 1997 DBS model. In those instances, the values that DBS used are recorded in the table. New parameters, or parameters for which the values used by DBS were not discernible are recorded with a "?" in that column. The values that we tested in the initial optimisation (the first 100 generations) are recorded in the next column. Some parameters are clamped in this phase, meaning that their values are not allowed to vary. After 100 generations, these parameters were allowed to vary. Some parameters were allowed to vary throughout the optimisation procedure, and were spawned in the first generation, meaning that values associated with the agents of the first generation were sampled from a normal distribution. The final column records the limits to which the parameter values were constrained throughout the optimisation procedure.

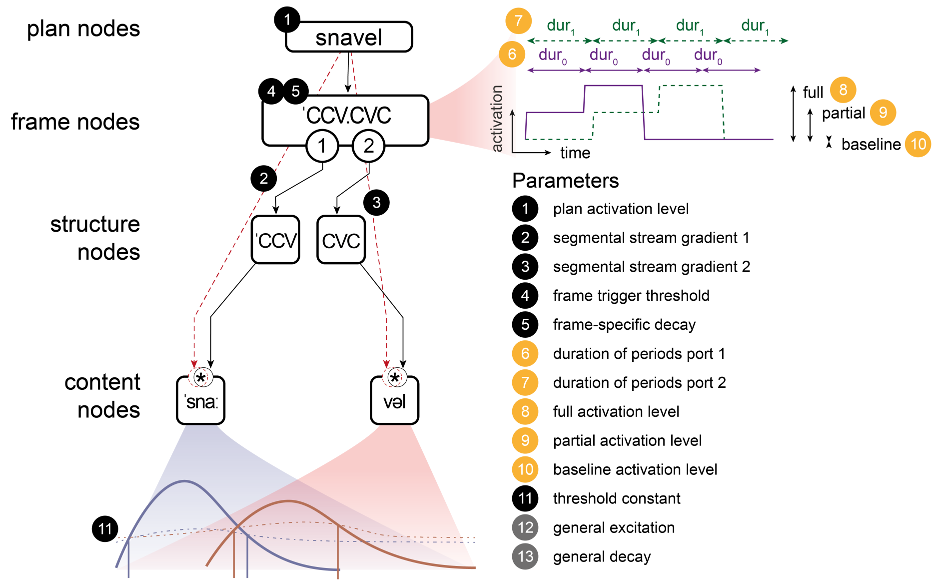

Figure 7: Parameters of the asynchronous model variant

Table 6: Parameters of the asynchronous model variant

	Parameter and description	value in DBS	Clamp in initial optimisation or spawn centre	Limits of parameter value
1	**plan activation** Activation level. This is the activation that is assigned to the plan node for the period from tick for to tick 28, during which time the plan node is constantly activated.	7	1.0, spawn centre	0.0 - Inf
2	**segmental stream gradient 1** Connection weighting. This controls what proportion of the activation in the plan node flows to the first content node.	Specified independently for each word	0.4, clamp	0.0 - 1.0
3	**segmental stream gradient 2** Connection weighting. This controls what proportion of the activation in the plan node flows to the first content node.	Specified independently for each word	0.7, clamp	0.0 - 1.0

Table 6: Parameters of the asynchronous model variant *(continued)*

	Parameter and description	value in DBS	Clamp in initial optimisation or spawn centre	Limits of parameter value
4	**frame trigger threshold** Connection weighting. This controls what proportion of the activation in the plan node flows to the first content node.	?	0.25, spawn centre	0.0 - 0.75
5	**frame-specific decay** Decay rate. The rate at which activation decays for the frame node.	0.4 - 0.6	0.9, clamp	0.0 - 1.0
6	**duration of periods port 1** (dur_0) Duration in ticks. Duration of each period of the activation pattern associated with the first port.	?	8, spawn centre	1 - 15
7	**duration of periods port 2** (dur_1) Duration in ticks. Duration of each period of the activation pattern associated with the second port.	?	8, spawn centre	1 - 15
8	**full activation** Activation level. The activation level of periods of full activation transmitted by the frame node ports.	1.0	1.0, spawn centre	0.0 - 1.0
9	**partial activation** Activation level. The activation level of periods of partial activation transmitted by the frame node ports.	0.4 - 0.6	0.5, spawn centre	0.0 - 1.0
10	**baseline activation** Activation level. The activation level of periods of partial activation transmitted by the frame node ports.	0.0	0.0, spawn centre	0.0 - 1.0
11	**threshold constant** The constant baseline of the threshold against which activation in the content nodes is compared to establish times of syllable onset and offset.	?	6.0, spawn centre	-6.0 - 18.0
12	**general excitation** Connection weighting. The weighting of all excitatory connections in the model that do not have a specific parameter defined weighting.	?	0.83, clamp	0.0 - 1.0
13	**general decay** Decay rate. The rate at which activation decays for all model nodes, except the frame node.	0.4 - 0.6	0.9, clamp	0.0 - 1.0

Supplement E: Multiple regimes to achieve the same speaking rate

In the first analyses of strand 2, we assumed that only one regime exists for each speaking rate. It is of course possible that several distinct combinations of parameter settings (regimes) might be able to account for the temporal structure of speech at each speaking rate; that is, there might be several different, equally viable "ways" to speak slow, fast, and at a medium rate. Such a situation would manifest itself in multiple distinct clusters in the parameter space emerging in the set of best performing unique solutions in each rate condition run. To establish whether multiple distinct regimes were present in the parameter values adopted to speak at each speaking rate, we performed k-means clustering on the parameter dimensions for the Pareto optimal solutions of each rate condition, which seeks to find k distinct clusters of points. We explored values of k from 1 to 39 (100 repetitions at each value of k), and calculated the Bayesian Information Criterion (BIC) for each clustering attempt. The BIC characterises the quality of an attempt, balancing the likelihood of each cluster with the number of clusters to avoid over-fitting. We retained the best clustering model for each value of k. For all three speaking rates, the best supported number of clusters was 1, meaning that the Pareto fronts contained only one "way" to achieve each speaking rate.

Supplement F: The loading of parameters onto the PCA

To be able to examine the arrangement in parameter space of the regimes associated with each of the rate conditions, we conducted a principal component analysis (PCA). This procedure loads as much variance as possible onto each component in turn, whilst ensuring that each component is orthogonal to the preceding PCs. PC1 (the first PC) accounted for 33.7% of the variance, PC2 accounted for 15.5% of the variance, PC3 for 11%, and PC4 for 6%. The loadings of the parameters onto the PCs are listed in Table 7.

Table 7: The loading of parameters onto principal components

	Parameter	PC1	PC2	PC3
1	plan activation (μ)	-0.164	0.488	-0.071
1	plan activation (σ)	-0.135	0.197	0.076
2	segmental stream gradient 1 (μ)	0.351	0.064	0.066
2	segmental stream gradient 1 (σ)	0.093	-0.133	-0.385
3	segmental stream gradient 2 (μ)	-0.312	-0.276	0.023
3	segmental stream gradient 2 (σ)	-0.164	-0.238	0.315
4	frame trigger threshold (μ)	0.141	-0.219	-0.268
4	frame trigger threshold (σ)	0.073	0.003	-0.143
5	frame-specific decay (μ)	0.078	0.03	0.508
5	frame-specific decay (σ)	0.005	-0.009	-0.277
6	dur_0 (μ)	0.246	-0.363	0.079
7	dur_1 (μ)	0.333	0.065	0.065
8	full activation (μ)	0.053	0.007	-0.064
8	full activation (σ)	0.248	0.183	0.188
9	partial activation (μ)	0.355	0.089	0.057
9	partial activation (σ)	-0.065	0.024	0.053
10	baseline activation (μ)	0.002	0.062	0.211
10	baseline activation (σ)	-0.019	0.04	0.15
11	threshold constant (μ)	0.337	-0.195	0.07
12	general excitation (μ)	-0.123	-0.532	0.052
12	general excitation (σ)	-0.084	-0.079	0.075
13	general decay (μ)	-0.328	0.074	-0.122
13	general decay (σ)	-0.142	-0.013	0.35
	Variance accounted for by principal component	33.7%	15.5%	11%

Figure 8 is an extended version of Figure 10 in the article text. The odd rows show the arrangement of the regimes on the planes of various pairs of principal components. Dots indicate members of the set of best performing solutions in each speaking rate condition (red for fast, green for medium and blue for slow). This visualises the arrangement of the parameters in PC space.

The plots in the even rows visualise the loadings of the parameters onto the PCs of the plane of the plot immediately above. The arrows point in the positive direction of the indicated parameter, the length of the arrow indicates the weighting. For the sake of readability, only parameters with an absolute loading greater than 0.35 on one of the PCs are included. Numbers refer to the parameter numbers given in the Table 7 and the table and in Figure 7 and Table 6 in Supplement D.

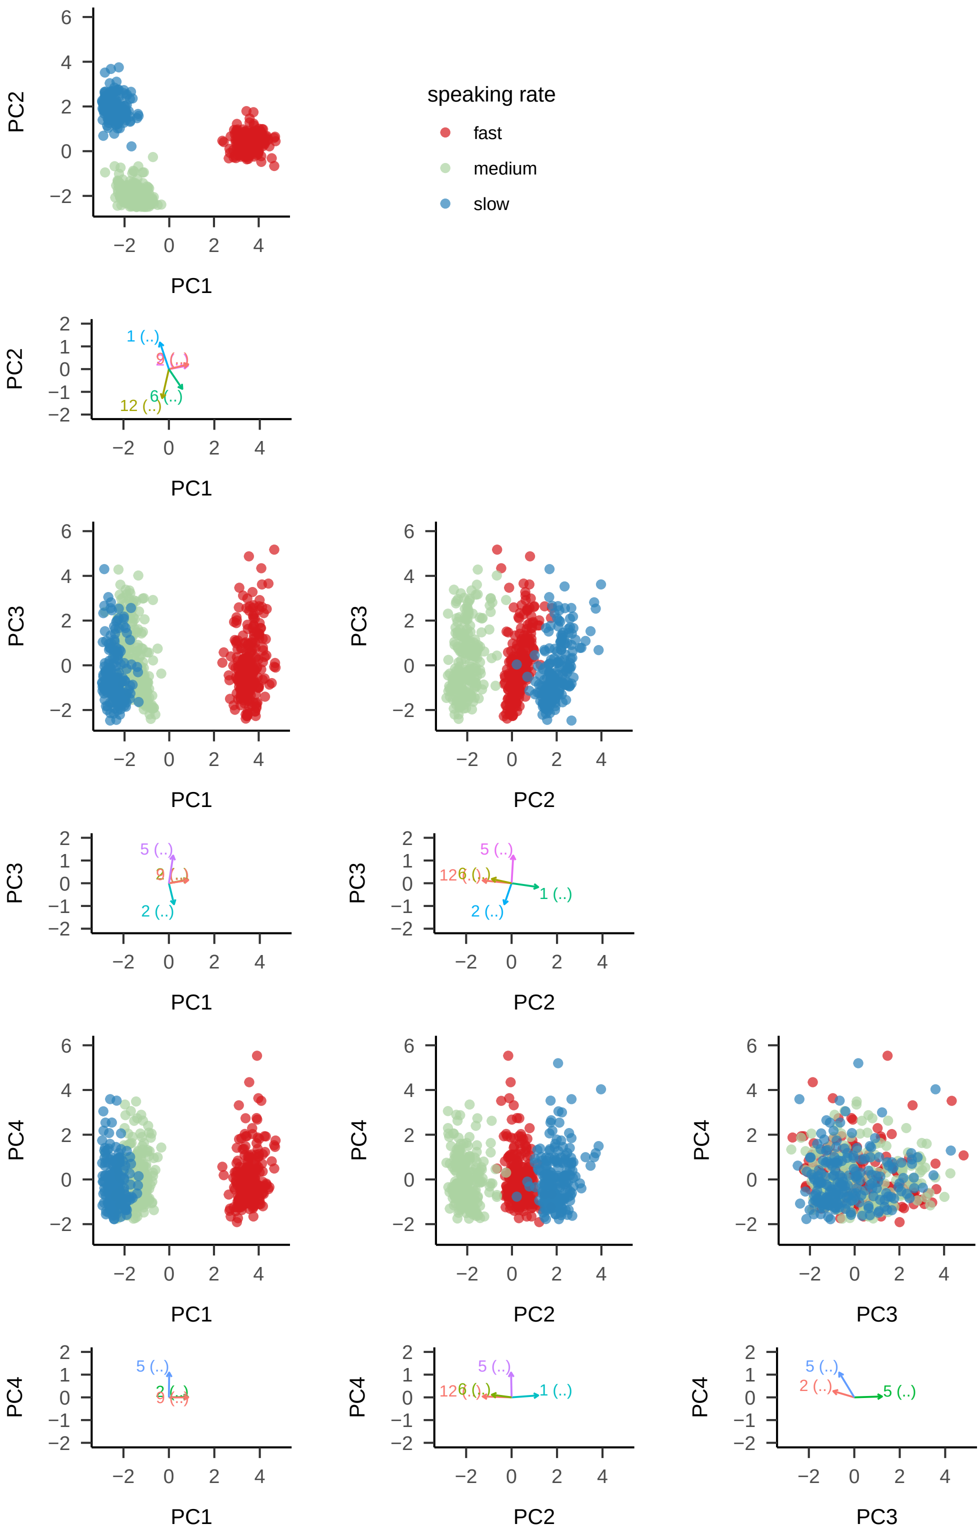

Figure 8: Odd rows: The Pareto optimal solutions identified for the fast (red), medium (green), and slow (blue) rate conditions, plotted for PC1, PC2 and PC3. Even rows: visualisation of the loading of parameters onto the PCs, see text for details.

Supplement G: Durations predicted by a sample of optimal solutions

Figure 9 shows predicted durations for a sample of solutions from the Pareto front, and is complementary to Figure 9 in the main text.

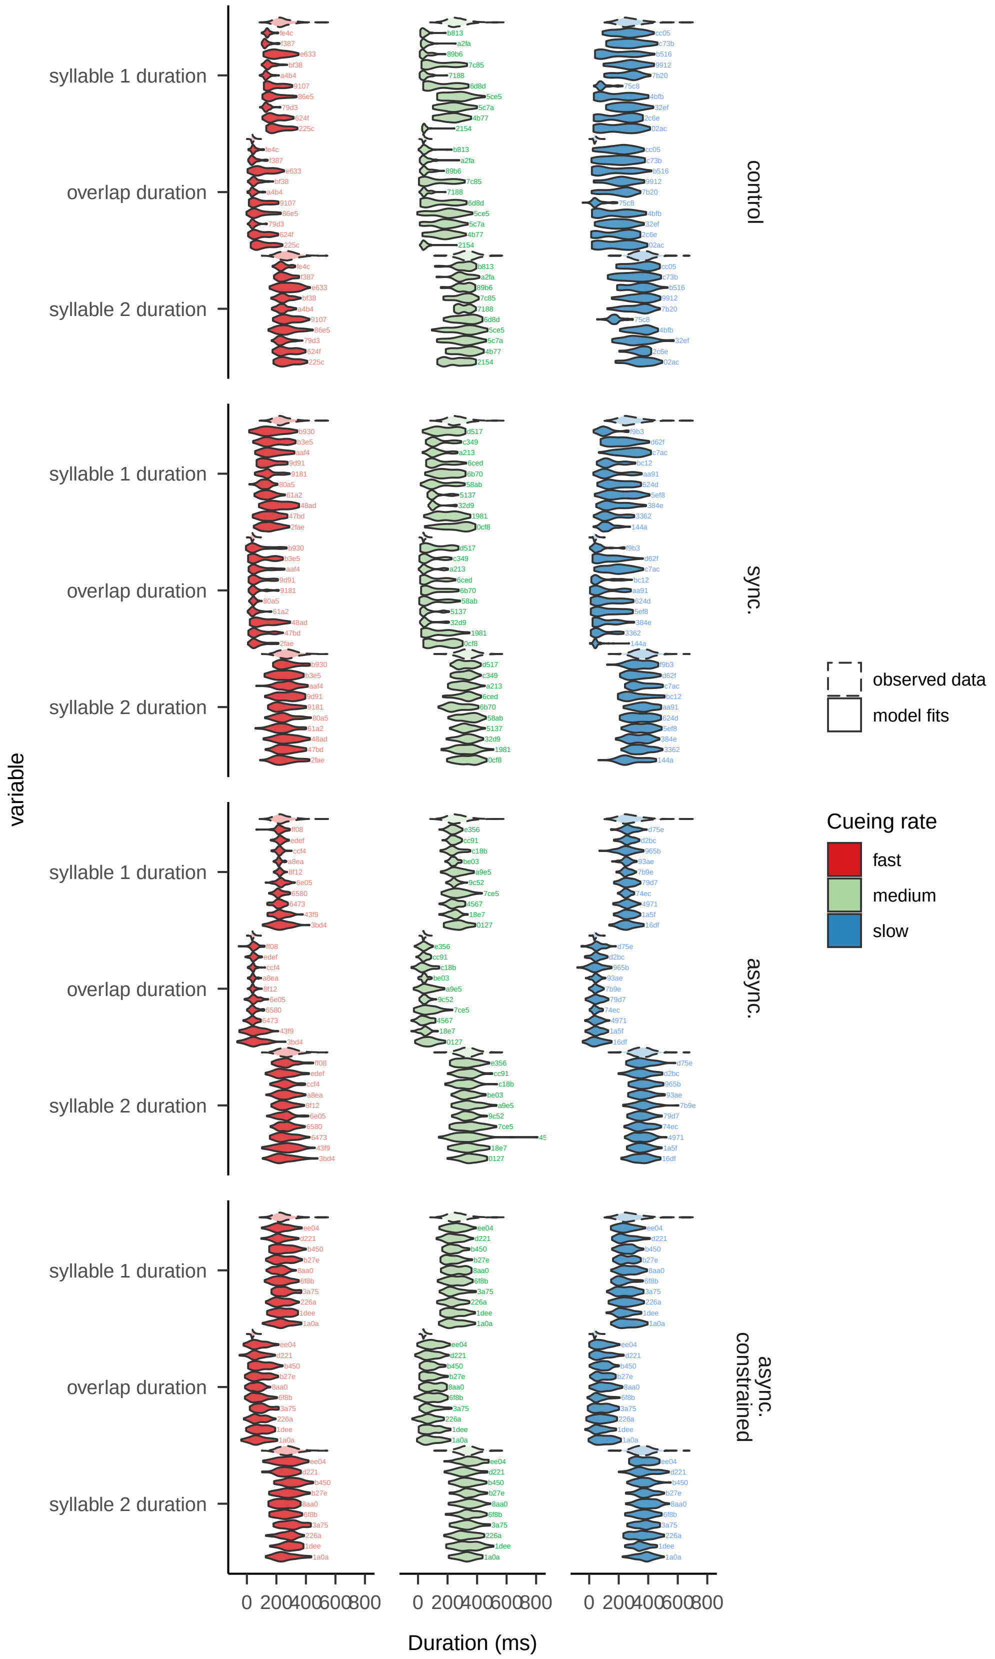

Figure 9: Predicted durations for a sample of solutions from the Pareto front

Supplement H: Details of GAM analysis of rate-to-gait mapping

For each axis of the extrapolated fingerprint duration data, we regressed the normalised fingerprint durations by the number of the step along the axis. Depending on the mapping to be modelled, this was either a "uniform" fit, or a "splined" fit. The uniform fit was a simple linear regression, whereby duration was predicted by the interaction between step number and a factor representing the three component fingerprint durations. The splined fit had a thin-plate spline smooth with a basis dimension of 4, grouped by the component factor. A basis of 4 means that the spline can adopt a sigmoid shape to fit the data.

We will take the example of the "fast is special" mapping, where fast is distinct and medium and slow are mapped to the same gait. The fit resulting from this model is depicted in the Figure 10. We construct a model consisting of a uniform fit for the axis between medium and slow, and splined fits for the axis between fast and medium and the axis between fast and slow. The multivariate nature of the model means that we can extract fit quality information for all three axes together, and thus evaluate statistical support for the five hypothesised mappings directly.

We constructed multivariate models fitting the three axes together, so we have one model for each hypothesis. The following tables summarise the regression coefficients of each the models in turn. Fitting was done using the *mgcv* R package (version 1.8-28) in R (R Core Team, 2018, version 3.5.2).

Results

Figure 11 presents the model comparison results of both the GAMs (in panel A) and the the Bayesian linear switchpoint models (in panel B). Panel B repeats the data shown in Panel E of Figure 11 in the article text. In both cases, we compare models on information criteria, which aim to quantify the explanatory power of the models in terms of the amount of information lost, while at the same time penalising model complexity to avoid over fitting. For the GAMs, we calculate the Akaike information criterion (AIC, Akaike, 1974), for the Bayesian linear switchpoint models we calculate an information criterion by leave-one-out cross validation (the LOOIC, Vehtari, Gelman, & Gabry, 2017).

The ordering of the mappings is the same according to both the GAM analysis and the Bayesian linear switchpoint analysis.

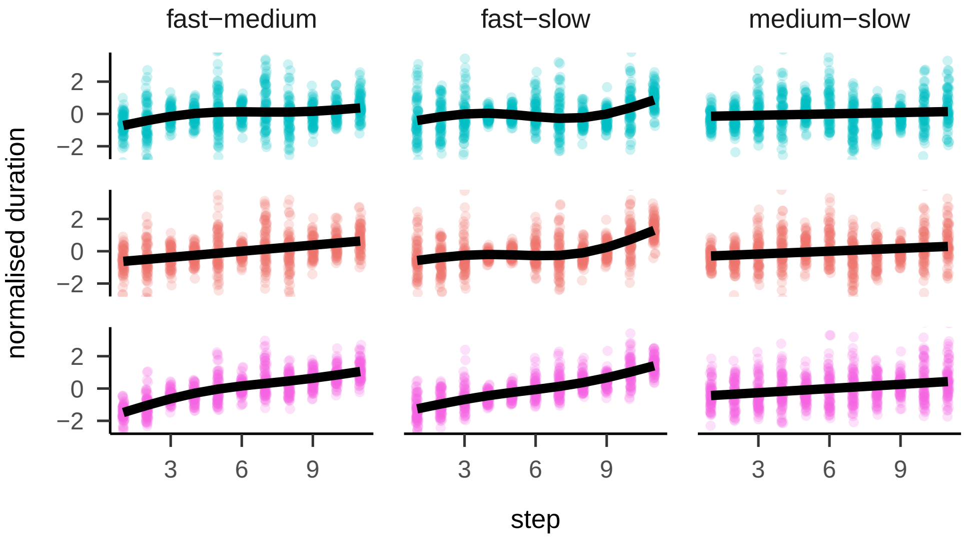

Figure 10: Example GAM fit for a 'fast is special' mapping

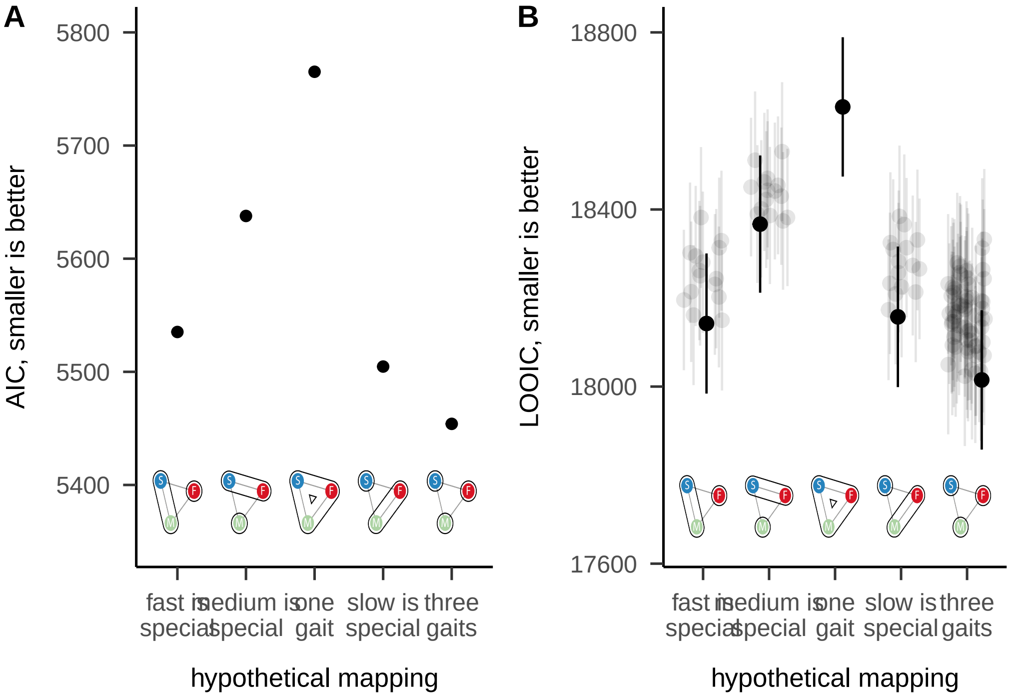

Figure 11: Comparison of information criteria for GAMs and Bayesian linear switchpoint models

Model of "fast-is-special" mapping

```
durfs ~ s(stepnr, by = fcomponent, k = 4)
durfm ~ s(stepnr, by = fcomponent, k = 4)
durms ~ stepnr * component
```

Linear coefficients

Axis	Parameter	Estimate	Standard Error	z value	p-value ($> \|z\|$)
fast-slow	(Intercept)	0.00	0.02	0.05	0.96
fast-medium	(Intercept)	0.00	0.02	0.22	0.82
medium-slow	(Intercept)	-0.17	0.07	-2.31	0.02
medium-slow	step number	0.03	0.01	2.64	0.01
medium-slow	syllable 1	-0.18	0.10	-1.75	0.08
medium-slow	syllable 2	-0.35	0.10	-3.36	0.00
medium-slow	step number:syllable 1	0.03	0.02	1.98	0.05
medium-slow	step number:syllable 2	0.06	0.02	3.78	0.00

Smoothing coefficients

Axis	Parameter	edf	Ref.df	χ^2	p-value
fast-slow	componentoverlapdur	2.96	3.00	140.66	0
fast-slow	smooth(stepnr):syllable 1	2.95	3.00	341.11	0
fast-slow	smooth(stepnr):syllable 2	2.74	2.95	765.16	0
fast-medium	componentoverlapdur	2.84	2.98	104.01	0
fast-medium	smooth(stepnr):syllable 1	1.00	1.00	178.09	0
fast-medium	smooth(stepnr):syllable 2	2.79	2.96	649.64	0

Model of "medium-is-special" mapping

```
durfs ~ stepnr * component
durfm ~ s(stepnr, by = fcomponent, k = 4)
durms ~ s(stepnr, by = fcomponent, k = 4)
```

Linear coefficients

Axis	Parameter	Estimate	Standard Error	z value	p-value ($> \|z\|$)
fast-slow	(Intercept)	-0.43	0.06	-6.70	0.00
fast-slow	step number	0.07	0.01	7.58	0.00
fast-slow	syllable 1	-0.42	0.09	-4.66	0.00
fast-slow	syllable 2	-1.06	0.09	-11.82	0.00
fast-slow	step number:syllable 1	0.07	0.01	5.26	0.00
fast-slow	step number:syllable 2	0.18	0.01	13.35	0.00
fast-medium	(Intercept)	0.00	0.02	0.22	0.82
medium-slow	(Intercept)	0.00	0.02	0.09	0.93

Smoothing coefficients

Axis	Parameter	edf	Ref.df	χ^2	p-value
fast-medium	componentoverlapdur	2.85	2.98	105.17	0
fast-medium	smooth(stepnr):syllable 1	1.00	1.01	178.09	0
fast-medium	smooth(stepnr):syllable 2	2.79	2.97	649.58	0
medium-slow	componentoverlapdur	2.95	3.00	47.18	0
medium-slow	smooth(stepnr):syllable 1	2.94	3.00	67.54	0
medium-slow	smooth(stepnr):syllable 2	1.46	1.77	70.53	0

Model of "one gait" mapping

```
durfs ~ stepnr * component
durfm ~ stepnr * component
durms ~ stepnr * component
```

Linear coefficients

Axis	Parameter	Estimate	Standard Error	z value	p-value ($> \|z\|$)
fast-slow	(Intercept)	-0.43	0.06	-6.70	0.00
fast-slow	step number	0.07	0.01	7.58	0.00
fast-slow	syllable 1	-0.42	0.09	-4.66	0.00
fast-slow	syllable 2	-1.06	0.09	-11.82	0.00
fast-slow	step number:syllable 1	0.07	0.01	5.26	0.00
fast-slow	step number:syllable 2	0.18	0.01	13.35	0.00
fast-medium	(Intercept)	-0.50	0.06	-7.69	0.00
fast-medium	step number	0.08	0.01	8.74	0.00
fast-medium	syllable 1	-0.26	0.09	-2.80	0.01
fast-medium	syllable 2	-0.91	0.09	-9.93	0.00

(continued)

Axis	Parameter	Estimate	Standard Error	z value	p-value ($>|z|$)
fast-medium	step number:syllable 1	0.04	0.01	3.16	0.00
fast-medium	step number:syllable 2	0.15	0.01	11.24	0.00
medium-slow	(Intercept)	-0.17	0.07	-2.31	0.02
medium-slow	step number	0.03	0.01	2.64	0.01
medium-slow	syllable 1	-0.18	0.10	-1.75	0.08
medium-slow	syllable 2	-0.35	0.10	-3.36	0.00
medium-slow	step number:syllable 1	0.03	0.02	1.98	0.05
medium-slow	step number:syllable 2	0.06	0.02	3.78	0.00

For the one gait model, there are no smoothing fits, since all three axes are fitted by uniform regressions only.

Model of "slow-is-special" mapping

```
durfs ~ s(stepnr, by = fcomponent, k = 4)
durfm ~ stepnr * component
durms ~ s(stepnr, by = fcomponent, k = 4)
```

Linear coefficients

Axis	Parameter	Estimate	Standard Error	z value	p-value ($>|z|$)
fast-slow	(Intercept)	0.00	0.02	0.05	0.96
fast-medium	(Intercept)	-0.50	0.06	-7.69	0.00
fast-medium	step number	0.08	0.01	8.74	0.00
fast-medium	syllable 1	-0.26	0.09	-2.80	0.01
fast-medium	syllable 2	-0.91	0.09	-9.93	0.00
fast-medium	step number:syllable 1	0.04	0.01	3.16	0.00
fast-medium	step number:syllable 2	0.15	0.01	11.24	0.00
medium-slow	(Intercept)	0.00	0.02	0.09	0.93

Smoothing coefficients

Axis	Parameter	edf	Ref.df	χ^2	p-value
fast-slow	componentoverlapdur	2.96	3.00	140.81	0
fast-slow	smooth(stepnr):syllable 1	2.95	3.00	342.20	0
fast-slow	smooth(stepnr):syllable 2	2.75	2.95	765.29	0
medium-slow	componentoverlapdur	2.95	3.00	49.48	0
medium-slow	smooth(stepnr):syllable 1	2.95	3.00	71.12	0
medium-slow	smooth(stepnr):syllable 2	1.52	1.85	70.47	0

Model of three gaits mapping

```
durfs ~ s(stepnr, by = fcomponent, k = 4)
durfm ~ s(stepnr, by = fcomponent, k = 4)
durms ~ s(stepnr, by = fcomponent, k = 4)
```

Linear coefficients

Axis	Parameter	Estimate	Standard Error	z value	p-value ($>|z|$)
fast-slow	(Intercept)	0	0.02	0.05	0.96
fast-medium	(Intercept)	0	0.02	0.22	0.82
medium-slow	(Intercept)	0	0.02	0.09	0.93

Smoothing coefficients

Axis	Parameter	edf	Ref.df	χ^2	p-value
fast-slow	componentoverlapdur	2.96	3.00	141.85	0
fast-slow	smooth(stepnr):syllable 1	2.95	3.00	342.20	0
fast-slow	smooth(stepnr):syllable 2	2.74	2.95	765.19	0
fast-medium	componentoverlapdur	2.85	2.98	105.08	0
fast-medium	smooth(stepnr):syllable 1	1.00	1.00	178.10	0
fast-medium	smooth(stepnr):syllable 2	2.79	2.96	649.57	0
medium-slow	componentoverlapdur	2.95	3.00	49.96	0
medium-slow	smooth(stepnr):syllable 1	2.95	3.00	71.11	0
medium-slow	smooth(stepnr):syllable 2	1.47	1.78	70.55	0

Supplement I: Details of Bayesian analysis of rate-to-gait mapping

For technical reasons, Supplement I, describing Bayesian switchpoint models could not be hosted by the APA. It is available at https://osf.io/3mqgu/.

Supplement J: Model and analysis code

For technical reasons, Supplement J, including the code of the model and analysis code could not be hosted by the APA. It is available at https://osf.io/3mqgu/.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE transactions on automatic control*, *19*(6), 716–723.

Cleveland, W. S., & Devlin, S. J. (1988). Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American statistical association*, *83*(403), 596–610.

Dell, G. S., Burger, L. K., & Svec, W. R. (1997). Language production and serial order: A functional analysis and a model. *Psychological review*, *104*(1), 123–147.

R Core Team. (2018). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria. Retrieved from https://www.R-project.org/

Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, *27*(5), 1413–1432. doi:10.1007/s11222-016-9696-4

Zitzler, E., Brockhoff, D., & Thiele, L. (2007). The Hypervolume Indicator Revisited: On the Design of Pareto-compliant Indicators Via Weighted Integration. In S. Obayashi, K. Deb, C. Poloni, T. Hiroyasu, & T. Murata (Eds.), *Evolutionary Multi-Criterion Optimization* (pp. 862–876). Springer Berlin Heidelberg.