

Supplementary Material


DYNAMICS-WEB: Web Content Character over the Link Graph

Jason Duke, Kronaxis Limited


0.1 10. The model on a real person, computed in the open

Everything to this point is either a measurement of content or a calibrated simulation of readers. The decisive question is whether any of it describes a real human, and the honest way to answer it is on a real person's own behaviour, with the workings open to inspection. This section is a single subject proof of concept and is offered as exactly that, a worked illustration of the method on one consenting person, not evidence about people in general; the cohort study that would turn it into a result, thirty or more participants with pre registered hypotheses and held out behavioural prediction, is set out in section 12 and is not yet run.


We took one consenting person, the author, and their actual browsing history, and asked whether their personality predicts the character of the content they really read. The first attempt nulled, and the null was instructive. Scored at the domain level, averaging the character of every page on each site visited, the profile was indistinguishable from the web average, because large diverse sites carry a blended, mean reverting character. The signal is not which sites you visit but which pages.



Scored at the page level the profile was clear and coherent: this reader takes deep, candid, authoritative and formal content and avoids emotional and commercial pages, the two largest departures from the web average. Then the test that matters. A DYNAMICS-8 profile inferred blind from behaviour alone, with the content profile withheld, predicted six of the seven non flat axes correctly under a rule fixed in advance, including the acuity to depth link the cross topic work had already singled out as the one that travels. The same prediction from a questionnaire scored four of seven. One person is a proof of concept, not a proof, but it is the line between a model that describes simulated agents and one that describes a human being.



Measuring the same person three ways then exposed a structure we were not looking for. Read as a natural self report, as a disciplined self report, and off behaviour, the cognitive traits, acuity and novelty, were near identical across all three: you cannot fake analytical sharpness and you do not need to. The self regulation traits, discipline, impulsivity and mercuriality, formed a clean gradient, the natural self impulsive, the enforced self dialled toward control, behaviour furthest toward control. The gap is the compensation, the habits a naturally impulsive person builds to manage themselves, and it appears precisely and only on the traits where willpower can act.



Personality has a felt layer that a self report captures and an enacted layer that behaviour captures, and for anything that turns on what a person actually does, what they read, buy or click, the enacted self is the relevant one, and it is written in behaviour rather than in self report. That is why the behavioural inference beat the questionnaire.


This is also where the toolbar earns its privacy claim rather than asserting it, because the client code is open and can be read line by line. The character of a page is computed on the reader's own device, by the distilled model of section 3 running inside the extension; the page content is never sent anywhere to be scored. What leaves the machine, and only under a consent the reader can see, is two things: the eight number character of pages, which extends the shared model so the next reader is served instantly, and the ongoing stream of pages visited and actions taken, sent to a real endpoint on our own servers, which is the behavioural record that turns the

single subject above into a population. The consent is on by default and the reader can switch it off at any time; because the client is open, anyone can verify that switching it off stops the transmission and that page content never leaves the device in either case. The per person signal lives at the level of the specific pages a person reads, a granularity no incumbent models per user, and that is both where the defensibility sits and where the responsibility does.

0.2 11. From one person to a population

The result of section 10 is a method, not a one off. The toolbar builds the same profile for any reader who wants one, by the same page level route, and offers three ways to do it. In live accrual, every page you browse is scored on your device as you go and the profile fills in over time. On device build reads your existing history, fetches and scores each page locally with the same model, and asks nothing of our servers but the result. Fast build scores your history on our servers when you would rather not wait. All three rest on a single primitive: a page's DYNAMICS-WEB character, computed wherever, tied to an anonymous per person key, from which we aggregate the reader's own DYNAMICS-8. The two volatility traits, Impulsivity and Mercuriality, are left blank by construction, because content cannot read them, which is the section 6 boundary showing up honestly in the product.

One person to a population: three ways to build a reader's profile



Page content is computed on your device and never uploaded. The character scores and your profile are sent, with consent, tied to an anonymous key.

The point of the productised form is the path it opens. The single subject of section 10 becomes a population the moment the same instrument runs for many readers, each contributing their page level character under consent. Page content stays on the device; the character scores and the finished profile are what travel, tied to a key that is not a name. The honest claim the paper can make today rests on one person; the honest claim it is built to make next rests on the many, and the toolbar is the instrument that gets from one to the other without ever shipping the mapping that turns a reader into the content that fits them.

0.2.1 11.1 The questions this opens

Once character is measured, propagated and paired to a reader, three empirical questions become askable that the field has not had the instruments to pose. None is answered here. Each is stated as a falsifiable hypothesis with a measurable outcome, in the same spirit as the human validation of section 5.4, so the next phase of work can confirm or refute it rather than assert it.

1. **Conversion, now answered for its first stage.** Section 5.6 reports the pre registered, powered experiment: tailoring the sub perceptual axes of a page to a reader's disposition

raised conversion from nineteen to thirty one per cent, a confirmed effect at a one sided p of 0.0295. What remains open is the crossover at scale, that each reader converts best on the character matched to them specifically rather than on a globally stronger page, which a larger pre registered Stage 2 is designed to confirm, and the replication of the whole effect across many readers and live commercial settings.

2. **Perceived quality.** Does the same tailoring raise the reader's own assessment of the page's quality, holding the information constant? This separates fit from fact. A page can carry identical content and be judged better or worse purely by how its character meets the reader, and the hypothesis is that character matched to disposition reads as higher quality to that specific person.
3. **Worldwide search.** Does the pairing generalise from a single page to ranking? A search engine that orders results by the fit between a searcher's inferred disposition and each result's character, the demand side principle set out in the companion Kronaxis Search discussion paper, should beat flat relevance for that searcher. Because section 5.3 shows the character construct travels across language families and scripts, the further question is whether the same character matched ranking holds worldwide rather than only in English. This is the largest of the three, and the propagated, multilingual character signal of this paper is built to make it testable.

Taken together they ask whether character, once measured, is not merely real and distinct but useful: whether matching it to a person changes what that person does, what they think of what they read, and what a search engine should put in front of them, anywhere in the world.

0.3 13. Prior art, owed plainly

This work meets several lineages and owes all of them plainly. Almost every idea it rests on belongs to someone else; the contribution is a specific combination they did not make, and one piece of it is the present author's own earlier work. We set the debts out in full rather than imply the field began here.

0.3.1 13.1 The graph and the text

On the graph side, topic sensitive PageRank (Haveliwala, 2002) and TrustRank (Gyongyi, Garcia Molina and Pedersen, 2004) established propagating a topical or trust prior through links, and Topical Trust Flow (Majestic) is the shipping commercial realisation of that idea at scale; our topical layer reimplements it on our own graph with no novelty claim over Majestic. Propagating character correctly also requires that affiliated neighbours not count as independent endorsements, the principle Hilltop introduced (Bharat and Mihaila, 2001), with the strength of a shared affiliation discounted by its rarity, the inverse log weighting of the Adamic-Adar index (Adamic and Adar, 2003). The particular synthesis the propagation uses, extending affiliation beyond domain and IP to registration and content and grading it by the rarity of the shared ownership cluster, the present author set out publicly in 2004 and an independent archive captured that year; the accompanying prior art note records the lineage and its dated source. On the text side the eight axes are not invented from nothing. Register, formality and stance are the substance of Biber's multidimensional analysis of register variation; affect is the province of sentiment and emotion classification; stance detection is an established task; and reading style off text is stylometry and authorship analysis. Several of our axes are, fairly stated, those constructs measured together on one scale.

0.3.2 13.2 The character of communication, and why it reads below conscious agreement

The idea that a piece of communication carries a measurable personality is not ours. Aaker (1997) built a validated five dimension scale of brand personality from human ratings, and the practice of scoring the tone of a message computationally is older still. Our character is the same kind of construct, moved from the brand to the individual page and from a survey to a machine. That human readers should agree on it only weakly, and feel it more than they can name it, is exactly what the processing fluency literature predicts: Reber, Schwarz and Winkielman (2004) show aesthetic and evaluative responses track the ease of processing, are hedonically marked, and are shifted by primes people cannot report, while Lindgaard and colleagues (2006) and later work show website impressions form in as little as fifty milliseconds, far below deliberation. Our human study, in section 5.4, recovers this dissociation for content character and we credit that literature for the mechanism; the new part is only that the same sub perceptual signal proves machine extractable at scale.

0.3.3 13.3 The look as a measurable, trust bearing channel

That a page's visual character is computationally measurable, and that it drives first impressions and credibility, is owed to a specific body of work. Reinecke and colleagues (CHI 2013) built computational models of perceived visual complexity and colourfulness from website screenshots and predicted first impressions from them, the very features our own visual analysis finds carry the most signal. Fogg's Stanford work on prominence and interpretation found the visual design of a page to be the single most cited factor in how people judge its credibility. Reinecke and Bernstein further showed that adapting a page's look to a viewer, in their case to cultural visual preference, measurably improves usability and preference. We extend this from a single aesthetic or credibility verdict, and from cultural adaptation, to a multi axis character read at web scale and matched to individual disposition, and we report honestly in our own scaling work that the look to character signal, strong on small curated samples, is modest on the open web and concentrated on the commercial and affective axes.

0.3.4 13.4 Matching content to the person, and the conversion evidence

The commercial premise, that matching content to a person raises conversion, is established, and we cite it rather than claim it. In advertising it is a century old: Hopkins (1923) ran coupon coded split tests of headlines, offers and look against measured response. In digital marketing the Eisenbergs' persuasion architecture (mid 2000s) matched content and experience to buyer personality types to lift conversion, and the conversion rate optimisation field has since run vast numbers of live tests of colour, tone, imagery and call to action against real outcomes. In the academy, Matz, Kosinski, Nave and Stillwell (2017) reached over three and a half million people and found that matching a persuasive appeal to a person's psychological profile produced materially higher click and purchase rates, with related effects in Hirsh, Kang and Bodenhausen (2012). The effect of matching is not in serious doubt. What none of this provides is the supply side: a way to measure the character of any page, at web scale, so that the matching can run against the whole web rather than against a handful of hand built variants or a single site's own copy. That measurement is what this paper adds, and section 5.6 closes the loop on our own axes: in a pre registered, powered test, matching the sub perceptual half of our machine measured character to a reader lifted conversion from nineteen to thirty one per cent at a one sided p of 0.0295. We borrow the established matching principle and now show it holds for the specific signal only this instrument recovers.

0.3.5 13.5 The author's own prior art, and the demand side

One debt is to the author's own earlier work, and it sets the priority of the demand side plainly. UK and US filing US20110197136A1, Web Browser Instance Profiling (inventor J. Duke, applicant StrangeLogic Ltd, priority date 10 February 2010, since lapsed and therefore public), disclosed a method of profiling a web browser instance to determine a user's propensity to undertake actions and then tailoring the appearance, content and incentives of a page to that inferred propensity. That is the demand side of the present thesis, profile the reader, infer disposition, match the content, set down in 2010, ahead of the academic confirmations cited above. The techniques it used belong to their time and have since been closed by browsers for privacy; the enduring idea, content tailored to an inferred disposition read from the reader's own browser, is recorded there. The mechanism that filing relied on, reading state straight from the browser, has since been closed for privacy, but the computation it reached for survives by a different route: a reader's disposition is recovered now from the character of the pages they actually consume, the method section 10 demonstrates, so the demand side runs as person to persona to content, model the reader as a disposition and match that disposition to content. The present work supplies the half that filing did not: the measurement of what every page's character actually is.

0.3.6 13.6 What is new here

Set against all of the above, the contribution is narrow and specific, and we claim only it: the treatment of character as a single, propagated, bidirectional, graph scale property of web content; the three readings computed across the whole host graph; the measured non redundancy of that signal, which reduces neither to authority (R squared 0.008 to 0.030) nor to surface text statistics (0.256, three quarters unexplained); its link homophily, strong enough to propagate (per axis neighbour correlation around 0.4 against a null of zero); the honest dissociation between a machine signal that is reproducible and a human judgement that is not; and the joining of this supply side measurement to a demand side profile the author first set out in 2010. We do not claim to have discovered register, sentiment, brand personality, processing fluency, visual aesthetics, or that personalisation works. We claim the propagated character signal and the measurement that makes matching possible against the whole web.