

This is a repository copy of *Motivated Agents:Toward the Computational Modeling of Motivational Affordances*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/146771/>

Version: Published Version

Conference or Workshop Item:

Deterding, Christoph Sebastian orcid.org/0000-0003-0033-2104 (2019) Motivated Agents:Toward the Computational Modeling of Motivational Affordances. In: Computational Modeling in Human-Computer Interaction, 05 May 2019.

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Motivated Agents: Toward the Computational Modeling of Motivational Affordances

Sebastian Deterding

Digital Creativity Labs, University of York

YO105GE York, UK

Sebastian.deterding@york.ac.uk

ABSTRACT

Video games routinely use procedural content generation, player modelling, and other forms of computational interaction that provide a good starting point for *engaging* computational interfaces. However, across these practices, games model environment (game content) and actor (player type) separately, which is out of tune with both basic and applied research. The ecological construct of motivational affordances, formalized as *actor-environment system ratios*, provides a promising alternative that could also prove fruitful for computational interaction in general.

INTRODUCTION

No matter if learning, productivity, behaviour change, wellbeing, or leisure: many interactive systems are tasked with reliably motivating their users. In response, HCI has developed theories and design formalizations for motivating interactions. From early work by Malone and Carroll on, one guiding intuition in this field is that video games form the practical avant-garde such motivating interactions, as they are purpose-built to afford enjoyable experiences [1]. And indeed, the intuition also holds for *computational interaction*, “the use of algorithms and mathematical models to explain and enhance interaction” [2]. Under the headers game analytics, player modeling, and procedural content generation (PCG), the video games industry routinely employs the mathematical analysis of large-scale user data, data-driven modelling of users, and artificial intelligence techniques to automatically generate and adapt interfaces and content [3,4].

Arguably the fullest realization of computational interaction for motivation can be found in experience-driven PCG [5], such as generating optimally fun *Super Mario Bros.* levels. Here, a generator produces content items like levels. An assessor then evaluates the items based on a player model, which could be an a priori computational model of intrinsic motivation or a neural network trained on user data. The assessor selects the content items it evaluates to most fit the

KEYWORDS

Computational modeling; motivation; affordances; motivational affordances; curiosity; difficulty balancing; novelty balancing

desired user experience. Players' actual engagement then is tracked and may in turn serve as input to the generator and/or update the player model.

THE CHALLENGE: SPLITTING ENVIRONMENT AND ACTOR

PCG, player modeling, and game analytics are not only instructive for any motivation-focused computational interaction: they also showcase a particular (and problematic) conceptual split: they model game content and player, actor and environment separately. On the one side stands *gamification* [1], trying to isolate motivationally 'active ingredients' –game design elements that reliably evoke the same motivational and behavioral response, irrespective of user or context. For instance, it tries to establish whether leaderboards reliably afford goal-setting. In contrast, *player modeling* [6] models traits or 'player types' as aggregate, cross-situational preferences for certain stimuli or behaviors. For instance, it tries to establish whether there is a "killer" player type that prefers to exert dominance over others, which a leaderboard may allow to. The two strands predictably meet in work on personalization that attempts to select game design elements that optimally fit the preferences of a given player (e.g. [7]). Yet as stated, gamification, player modeling, and game personalization all ultimately treat content and user as separable constructs.

This separate modeling is arguably a conceptual heritage of classic cognitive science, which informs strong-going "paradigm 2" HCI [8]. Yet influential as it has been, the classic cognitive science story struggles to account for large bodies of empirical research that underwrite contemporary ecological, enactive, embodied accounts. These accounts see actor and environment not as sharply separate, but enmeshed in a continuous perception-action loop, where actors actively move through environmental media to reveal directly action-relevant information that comprises environmental and bodily states in their relation [9].

Applied research concurrently suggests that decontextualized and on their own, neither environmental features (game design elements) nor user dispositions (player types) reliably predict motivational outcomes. Thus, we found that one and the same design element - a badge - would manifest different motivational functions and behavior depending on how users appraised them. This appraisal was afforded but not determined by the specifics of both usage context and design element [10]. Our work on contextual autonomy support in video game play showed that one and the same design and context feature could have diametrically opposed motivational outcomes depending on how it related to context and current user dispositions. To wit, a branching narrative could afford autonomy in leisurely play but thwart autonomy if part of a work task that involved methodically documenting the whole game. Pre-planned play time would only be perceived as autonomy-thwarting if the demanded play time window mismatched the player's own spontaneous wants and situational opportunities [11].

THE ALTERNATIVE: MOTIVATIONAL AFFORDANCES

One construct that spans the environment-actor divide is *affordances*. First developed by ecological psychologist J.J. Gibson, affordances capture directly perceivable actionable properties of the environment to an organism, such as the "sit-on-ability" of a chair (see [20] for a review).




Figure 1: Modelling affordances as actor-environment ratios. Warren [14] modelled climb-ability as the ratio of leg length ($L+L_1-L_2$) to stair Riser height R , $\pi = R/L$. Individuals converge around $\pi=0.24$ as the materially least energy-consuming and perceptually most 'climbable' ratio, and $\pi=0.88$ as the maximum climbable ratio. Figure taken from [14].




Figure 2: Flow or "optimal experience" as a ratio of actor skill S and environmental challenge C , $\pi = C/S$, suggesting an optimum $\pi=1$, with $\pi>1$ causing anxiety and $\pi<1$ causing boredom. Figure adapted from [15].

Importantly, affordances are conceived to exist only *in relation to* particular organisms: grass is and appears as edible to a cow, not a whale. It is simply meaningless to speak of the edibility or sit-on-ability of an object without reference to a particular organism. Affordances are central to contemporary enactive, embodied, ecological accounts of action and perception, and have been actively adopted in HCI and design, albeit often parsed in cognitivist, paradigm 2 terms [12].

Taking a cue from this, researchers have suggested *motivational affordances* as a construct to capture the motivating qualities of interactive systems (again often in cognitivist terms) [13]. In previous work, I tried to articulate a fully ecological, relational conception of "situated motivational affordances" as "the opportunities to satisfy motivational needs provided by the relation between the features of an artifact and the abilities of a subject in a given situation" [13]. This conceptualization accounts for aforementioned findings that humans encounter a rich socio-cultural world where objects have different uses and meanings depending on situational context.

While the concept has found adoption and inspired design work, it does not readily indicate ways of formalizing motivational affordances, which is crucial for computational interaction. Here, ecological psychology can provide inspiration. Warren [14] notably operationalized affordances as *system ratios of environmental to actor variables*. He successfully predicted the perceived and actual climb-ability of stairs as a ratio of stair step height to an individual's leg riser height: the metabolically optimal, least energy-consuming ratio was also freely chosen by participants on pure sight as the most desirable one (figure 1).

This move to actor-environment system ratios may not be always readily apparent. (One refrain of ecological psychology has been that identifying the relevant actor and environment properties for the affordances of most tasks remains an open empirical task.) But logic suggests that where these are identified, an affordance-as-system-ratio model should be more predictive than modeling either side individually. To take the legs and stairs example: If we were to model climb-ability as a function of stair height alone, we would get people's varying leg heights as noisy deviation around an aggregate mid-point. Similarly so if we were to model climb-ability as a function of leg height against the naturally occurring variance of step heights.

To illustrate the practical applicability of this approach to motivational affordances within a broader project of computational interaction, I will give two examples of currently ongoing projects, one on difficulty balancing and one on novelty balancing.

CASE 1: DIFFICULTY BALANCING

Difficulty balancing is a well-established practice in game design [16], grounded largely in Csikszentmihalyi's flow model [15] (figure 2). Following Csikszentmihalyi, activities tend to engage when their difficulty matches our skill. If difficulty is lower than our skill, the task is boringly easy. If difficulty exceeds our skill, the task induces anxiety and frustration. As people's skill tends to improve with practice, task difficulty needs to increase in lockstep to keep people engaged. Informed by Csikszentmihalyi, game designers try sequence game tasks in a well-formed difficulty curve. Yet since players' skills and learning speeds differ (one curve doesn't fit all), designers use various means of difficulty adjustment [16], ranging from player options (e.g. selecting between




Figure 3: In the crowdsourcing game *Paradox*, players try to manipulate nodes in a visual network to reach a desired 100% rate of satisfied edge requirements, covertly doing a software verification task [17]. We implemented the player rating system Glicko-2 to match players and crowdsourcing tasks/levels, using player winning odds E_p as our target affordance, which can be calculated from the relation of player and level rating, $E_p(r, v) = 1/(1 + 10^{(v-r)/400})$, where r is the player and v is the level rating. Figure from [17].

easy, medium, and hard mode) to systems automatically adapting difficulty in response to player performance.

In prior work, we extended such dynamic difficulty adjustment to crowd work, specifically the crowdsourcing game *Paradox* [17] (figure 3). Crowdsourcing games and other crowdsourcing platforms are known to suffer from poor retention: the majority of volunteers leaves after a short engagement, never to return. One reason is that crowdsourcing platforms rarely provide any difficulty balancing. Tasks are served either at random, or to optimize informational gain on the task, not user engagement. As a result, users may be early on served far too difficult tasks, triggering frustration and abandonment [17].

Unlike game content, crowdsourcing tasks pose the problem that tasks aren't created or manipulated at will (they depend on the job to be done), nor is their individual difficulty known in advance. To circumvent this issue, we repurposed player ranking systems to select (not create or adapt) tasks with 'fitting' difficulty. Roughly, ranking systems like the ELO Chess ranking give each player a numerical rank, where the delta between two matched players can be used to calculate the winning odds of the match. The actual match outcome is then used to update the ranking score of each player. ELO derivatives like Glicko or Trueskill combine this with confidence ratings based on how many matches have informed each player rank. In our system, we treated tasks as players, banking on the fact that each task is usually solved by more than one player for verification purposes. Thus, we have a uniform variable – a rank – that expresses both player skill and task difficulty. The relation of both – the rank delta – allowed us to predict the difficulty level of each user-task match. We experimentally compared a system randomly serving tasks to one serving tasks in uniform increasing difficulty and one matching difficulty to skill, and found that the matchmaking system led users to solve more, and more high-difficulty tasks [17].

CASE 2: NOVELTY BALANCING

Curiosity has been long considered an important motive driving exploratory behavior, play, and learning. Early work established a range of stimulus features that stoke curiosity, such as novelty, complexity, uncertainty, or conflict [18]. As with flow, there is some of an optimal 'goldilocks' mid-point for complexity and novelty [21]. Evidence also suggests that the effect of novelty and complexity charts a *relation* to the individual's knowledge. People more knowledgeable in music or painting would voluntarily expose themselves longer to more complex visuals or melodies [18].

Current work in cognitive science is converging around the idea that curiosity is a crucial motive driving organisms to explore their environment to increase their ability to predict future internal and external states [9]. Actors aim to predict and realize desirable future states, which entails orienting themselves towards and collecting more input on instances where their prediction mismatches perception. The higher the prediction error or surprisal, the more unknown, uncertain, and curiosity-inducing a phenomenon is – which again traces a relation between actor (prediction) and environment (perceived reality). The optimal learning strategy is to actively seek out instances where the organism expects the steepest decline in prediction error [19].

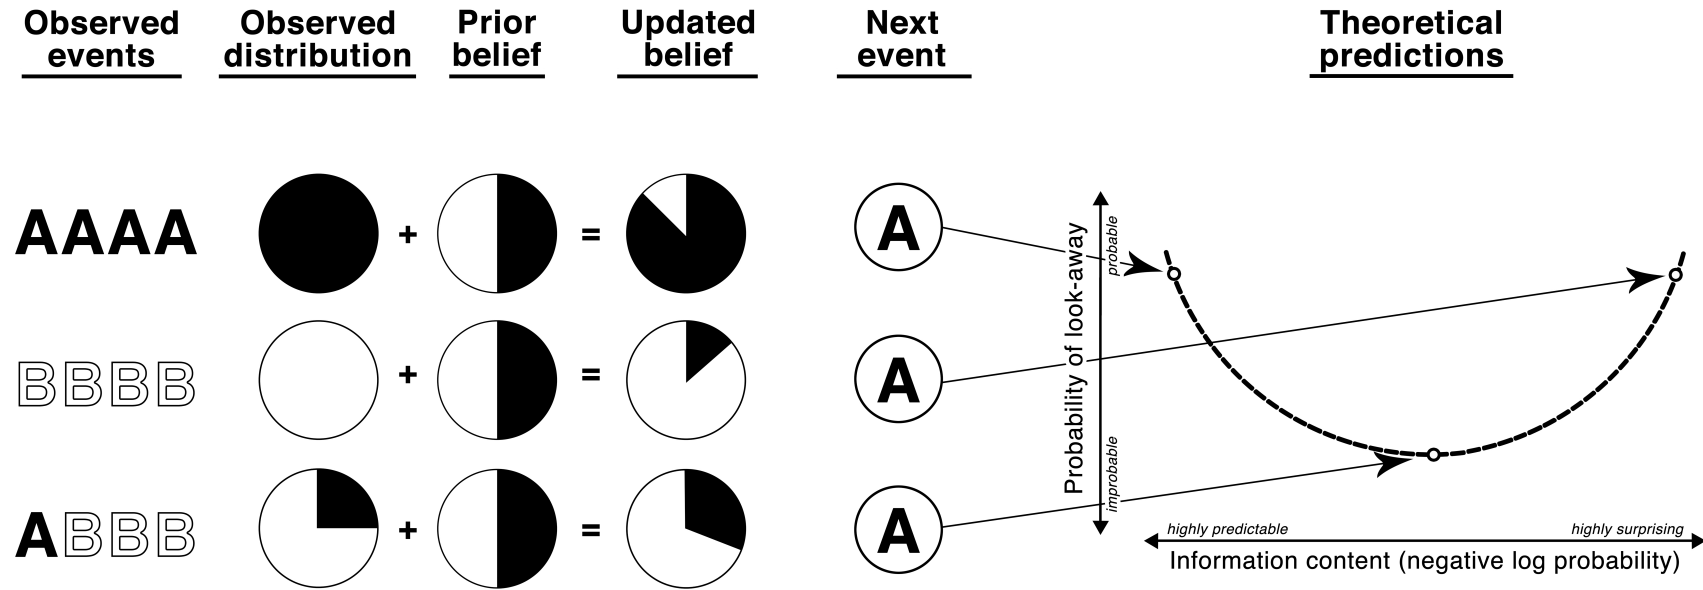


Figure 4: Kidd et al. [21] use a Markov Dirichlet-multinomial model to implement and test a simple Bayesian ideal observer model, predicting that the likelihood of looking away (as a behavioural expression of lost interest) is predicted by the expected (negative log) probability of the next presented stimulus. Empirical data with 7-8m old infants confirmed their model prediction of a ‘goldilocks’ midpoint. Again, the motivational affordance ‘interestingness’ (information content) is modelled as a relation of actor (prior belief) and environment (next event). Figure taken from [21].

Conveniently, this model has been formalized in e.g. “active sampling” learning algorithms for artificial agents [19] and computational models of e.g. attention [21] (figure 4). In current work, we are trying to transfer these formal models of curiosity into computational interactions for *novelty balancing*. We plan to use PCG frameworks to generate *Super Mario Bros.* level sequences that are then played by artificial agents which will provide novelty ratings for each level and the overall sequence as an aggregate of prediction errors experienced across each level. We will then compare this rating with human player ratings of novelty and curiosity to see whether their experiences track the computational measure.

CONCLUSION

Video games feature rich precedents that can inspire computational interaction for motivational interfaces. Yet like computational interaction more broadly, games tend to model actor and environment as separate, which is contradicted by current cognitive science and produces unreliable systems. In this paper, I argued for affordances, modelled as actor-environment system ratios, as a possible alternative. I illustrated this alternative approach with two projects. While they leave many questions unanswered and may turn out impractical for many use cases, I believe they still give an important impulse to advance the wider discourse around computational interaction.

REFERENCES

- [1] Sebastian Deterding. 2015. The Lens of Intrinsic Skill Atoms: A Method for Gameful Design. *Human-Computer Interaction* 30, 3–4: 294–335. <https://doi.org/10.1080/07370024.2014.993471>
- [2] Xiaojun Bi et al. 2018. Introduction. In *Computational Interaction*. Antii Oulasvirta et al. (eds). Oxford University Press, Oxford, 1–11.
- [3] Magy Seif El-Nasr, Anders Drachen, and Alessandro Canossa. 2013. *Game Analytics*. Springer, London. <https://doi.org/10.1007/978-1-4471-4769-5>
- [4] Noor Shaker, Julian Togelius, and Mark J. Nelson. 2016. *Procedural Content Generation in Games*. Springer, Cham. <https://doi.org/10.1007/978-3-319-42716-4>
- [5] Georgios N. Yannakakis and Julian Togelius. 2011. Experience-driven procedural content generation. *IEEE Transactions on Affective Computing* 2, 3: 147–161. <https://doi.org/10.1109/T-AFFC.2011.6>
- [6] Georgios N. Yannakakis et al. 2013. Player modeling. Dagstuhl Follow-Ups. Vol. 6. Dagstuhl-Leibniz-Zentrum fuer Informatik, 2013.
- [7] Gustavo F. Tondello, Rita Orji, and Lennart E. Nacke. 2017. Recommender systems for personalized gamification. In *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization*.
- [8] Steve Harrison, Deborah Tatar, and Phoebe Sengers. 2007. The Three Paradigms of HCI. In *ALT.CHI 2007*.
- [9] Andy Clark. *Surfing Uncertainty. Prediction, Action and the Embodied Mind*. Oxford, Oxford University Press.
- [10] Rob van Roy, Sebastian Deterding, and Bieke Zaman. 2018. Collecting Pokémon or receiving rewards? How people functionalise badges in gamified online learning environments in the wild. *International Journal of Human Computer Studies*, September 2018: 0–1. <https://doi.org/10.1016/j.ijhcs.2018.09.003>
- [11] Sebastian Deterding. 2016. Contextual Autonomy Support in Video Game Play. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*, 3931–3943. <https://doi.org/10.1145/2858036.2858395>
- [12] Shaleph O'Neill. 2008. Affordances: A Case of Confusion. In *Interactive Media: The Semiotics of Embodied Interaction*. Springer, London, 49–66. <https://doi.org/10.1007/978-1-84800-036-0>
- [13] Sebastian Deterding. 2011. Situated motivational affordances of game elements: A conceptual model. In *CHI 2011 Workshop "Gamification."* Retrieved from <http://gamification-research.org/chi2011/papers>
- [14] W H Warren. 1984. Perceiving affordances: visual guidance of stair climbing. *Journal of experimental psychology. Human perception and performance* 10, 5: 683–703.
- [15] Mihaly Csikszentmihalyi. 1990. *Flow: the psychology of optimal experience*. Harper and Row, New York.
- [16] David Michael and Jordan Chang. 2013. Dynamic Difficulty Adjustment in Computer Games. In *Proceedings of the Interactive Multimedia Conference*, 1–6.
- [17] Anurag Sarkar, Michael Williams, Sebastian Deterding, and Seth Cooper. 2017. Engagement Effects of Player Rating System-Based Matchmaking for Level Ordering in Human Computation Games. In *Proceedings of FDG'17*. <https://doi.org/10.1145/3102071.3102093>
- [18] Paul J. Silvia. 2006. *Exploring the Psychology of Interest*. Oxford, Oxford University Press.
- [19] Jacqueline Gottlieb and Pierre-Yves Oudeyer. 2018. Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience* 19: 758–770.
- [20] Anthony Chemero. 2009. *Radical Embodied Cognitive Science*. MIT Press, Cambridge, MA, London.
- [21] Celeste Kidd, Steven T. Piantadosi, and Richard N. Aslin. 2012. The Goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS ONE* 7, 5: 1–8. <https://doi.org/10.1371/journal.pone.0036399>