Citation: Sosa, R. & van Dijck, M. (2021) A Computational Interrogation of "Big-C" and "Little-c" Creativity, *Creativity Research Journal*, DOI: 10.1080/10400419.2021.1992195

# A Computational Interrogation of "Big-C" and "little-c" Creativity

The distinction between "Big-C" and "little-c" creativity implies that the generative process of celebrated creators is of a special type or degree. Arguments for and against such hierarchy of creativity are found in the literature, primarily built on rhetorical argumentation. The aim of this work is to examine the rationale behind Big-C and little-c creativity using explicit and more systematic means of inquiry. We employ computational agent-based simulations to study these constructs, their premises and their logical implications. The results of this work indicate that hierarchies such as the Big-C and little-c of creativity fail to provide a consistent way to explain and distinguish the generative processes of individual creators. In these computational models of creative social systems, only about half of disruptive changes can be explained by the characteristics of individual agents. This shows how labels like Big-C that are dependent on evaluation outcomes can easily be misattributed by observers to individual creators. This work demonstrates how the use of computational simulations can be useful to examine fundamental ideas about creativity. It shows that the Big-C/little-c distinction is a false dichotomy that should be approached critically by scholars to avoid conflating generative and evaluative dimensions of creativity.

Keywords: personal creativity; historic creativity; creative systems; creativity evaluation; creative systems; agent-based modelling

## Introduction

A distinction between "genius-level vs garden-variety" levels of creativity is often construed in the academic literature (Merrotsy, 2013, p. 474). The terms "Big-C" and "little-c" have helped cement such dichotomy by distinguishing between two alleged types of creative processes: those displayed by eminent creators and those found in everyday creators (Runco, 2014). This demarcation is used by scholars to articulate "ways that creativity has and should be conceptualized" (Kaufman & Beghetto, 2009, p. 2). A second and more formal definition hinges on an idea's probability, its utility, and its obviousness as estimated by an individual (little-c) or as agreed by a societal group (Big-C) (Simonton, 2017). This paper examines the former, more informal, and more prevailing use of the terms, since they remain largely unexamined and are used to think about creativity in ways that can carry important entailments.

The Big-C/little-c distinction has received considerable attention as evidenced by more than 2,000 citations to what was initially framed as a "preliminary, conceptual" model to "more clearly articulate the nature of creativity" (Kaufman & Beghetto, 2009, p. 1). Such hierarchy was formulated rhetorically and illustrated anecdotally (Kaufman & Beghetto, 2009) with Big-C described as the "remarkable and lasting contributions made by mavericks" (p. 2), and little-c portrayed as the "more incremental contributions made by everyday people" (p. 2). Thus, the core criterion to label a contribution as either Big-C or little-c does not refer to its creation, but depends on how it is judged in a social system (Csikszentmihalyi, 2014).

Extending this dichotomy of creative processes, a four-level hierarchy is further articulated by Kaufman and Beghetto (2009) to label the creativity displayed by young children rediscovering an art technique (first level: "mini-c") to the amateur artist who uniquely adapts a technique (second level: "little-c"), to the "highly accomplished but not yet eminent" professional (third level: "Pro-c") and, ultimately, to celebrities whose "works have lasted centuries" or who have won a prestigious award (fourth level: "Big-C") (Kaufman & Beghetto, 2009, p. 2). It can be useful to distinguish between two, four, or more levels of creative contributions "to separate personal from historical creativity" when it comes to the evaluation of ideas (Runco & Jaeger, 2012, p. 95). Others have further suggested that "even the four kinds of creativity need to be subdivided, because they are regions on a continuum, not four discrete kinds" (Sternberg, 2018, p. 4). Whether two or more, discrete or continuous, the distinction between levels of creativity is pervasive in the literature despite being problematic since it is considered to lack "real substance or validation" (Runco, 2014, p. 131).

A significant problem arises when the Big-C/little-c distinction is transferred from idea evaluation to idea generation. This is normally done by starting from the conclusion, namely: since Big-C ideas are deemed as "remarkable contributions" by a community (of experts, audiences, or consumers), it follows that *there must be* something special and unique about the creative process of celebrated creators (upper case Creativity) which is somehow necessarily distinct from the creative process of everyone else (lower case creativity). In the process of construing and contrasting these alleged levels of creativity, generative and evaluative factors are conflated whenever Big-C and little-c labels are applied based on events "that occur after the creative act" (Runco, 2014, p. 132). And since the Big-C and little-c labelling of an idea is only made after its social evaluation, it is problematic to use it to characterize the idea generation process much less the individual who generated the idea: "people who are very creative but not at the Big-C level are considered to be at the little-c level" (Kaufman & Beghetto, 2009, p. 3).

Beyond the discursive argumentation for or against Big-C/little-c distinctions and their usage, the aim of the work presented here is to more systematically and explicitly examine the rationale behind Big-C and little-c creativity. This examination is carried out by building and studying computational agent-based models to examine these constructs, their bases and their logical implications.

Agent-based models use computational simulation as a tool for empirical research to complement deductive and inductive approaches to social inquiry. A core idea of this "generative social science" (Epstein, 1999, p. 41) is that the phenomena of interest are *grown* (p. 43) in models that are exhaustively defined through algorithmic descriptions and act as testbeds for ideas and explanations (Calder et al., 2018, p. 2). As a modelling approach, these “artificial societies” are suitable for the study of heterogeneous agents with bounded rationality and whose autonomous behavior at the local level produce aggregate emergent structures that can shape back individual behavior. This modelling work consists of explicitly defining the rules for every agent, running the system over simulated time, and closely assessing the emergent aggregate patterns. An iterative process of construction and testing allows for (and demands) the type of conceptual clarity necessary to reason systematically about core premises such as the relationship between generation and evaluation of ideas in the definition of Big-C/little-c creativity. Through direct access to the source computer code of such models, issues of replicability and reliability are addressed.

Computational social simulations have a variety of applications including prediction, explanation, visualization, and theorization (Calder et al., 2018, p. 4). The present study has this last purpose and is thus used as an aide to formalize intuitions and thought experiments with the advantage that it offers a strict inferential process for rigorous reasoning and increased realism compared to mathematical models (Rangoni, 2014). We thus define, implement, analyze, and share an agent-based system here with the goal to test and check the Big-C/little-c creativity hypothesis.

In a typology of social simulation models (Gilbert & Ahrweiler, 2006), our current approach is closer to nomothetic science as our concern is to rigorously test a hypothesis, and further from idiographic science which uses data from the real world and is concerned with empirical validation (Gilbert & Ahrweiler, 2006, p. 22). In nomothetic agent simulations the intention is to formulate and test mechanisms or first-principles, whilst complementary idiographic modelling deals with specific instances and case studies (Salvatore & Valsiner, 2010, p. 829).

The main dimensions of concern in nomothetic agent-based modelling are related to their inbuilt biases and their usefulness to capture phenomena of interest (Carley, 2019, p. 743). In this work, we build and study a small-scale agent system (Sun et al., 2016) to bridge a priori and a posteriori modes of reasoning (Boghossian & Peacocke, 2000). This offers explicit and verifiable means to examine first-principles about Big-C and little-c creativity which are easier to describe and analyze than large-scale agent systems with multiple layers of causality and emergence.

Paraphrasing (Kitcher, 2000, p. 66), nomothetical agent simulations are employed for the inquiry of propositions *p* supported by the following account of a priori knowledge:

> *X* knows a priori that *p* iff *X* knows that *p* and *X*'s knowledge that *p* was produced by process α which is an a priori warrant for *p*. α is an a priori warrant for *p* just in case α is a process such that for any sequence of experiences sufficiently rich for *X* for *p*, some process of the same type could produce in *X* a belief that *p* (Kitcher, 2000, p. 66).

In this work, *p* stands for the Big-C/little-c distinction, the computational agent-based model provides the "sequence of experiences sufficiently rich", and "the same type" refers to the value of the model to aid researchers think about the structure, variables, and processes of *p* (Davis & O'Mahony, 2019, p. 29). This strategy allows to examine the dichotomy of little-c vs Big-C creativity beyond rhetorical arguments and in an artificial social system where we can grow and closely study the creation and evaluation of creative designs by social agents (Epstein, 1999). We thus use agent simulation as an "intuition pump" (Dennett, 2013) for exploratory analysis where "the purpose is not to predict what will happen, but to understand what *may* happen, and to estimate the circumstances under which various behaviors are most likely" (Davis & O'Mahony, 2019, p. 30).

Citation: Sosa, R. & van Dijck, M. (2021) A Computational Interrogation of "Big-C" and "Little-c" Creativity, *Creativity Research Journal*, DOI: 10.1080/10400419.2021.1992195

**An Agent-Based Model of Creative Design**

The advantages of agent simulation as an *intuition pump*, or a testbed for ideas, include overcoming human decision biases, thinking through a large number of alternative situations, breaking down causality, and reasoning over periods of time that are too complex for the human unaided by the computation to grasp (Carley, 2019, p. 741).

In this study, we model a social system of creative designers where new solutions are introduced by individual design agents and evaluated by a social group or a diverse community of design agents. To do this, our model captures a design task that goes beyond a search process in a solution space. It captures the significant invariants of design-task environments, namely: incomplete availability of information; negotiable constraints; interconnectivity of parts; no single right or wrong answers, only better and worse ones; and indirect and delayed feedback (Goel & Pirolli, 1992, p. 401).

The inspiration for the task captured in this agent simulation comes from two online communities where humans generate, share, and evaluate color palettes over many years: Adobe Color[1] and ColorLovers[2]. For a detailed description of those creative communities, refer to (O'Donovan, 2015, p. 19). In a nomothetic approach, we build this model using these communities as a reference, rather than as input data which would be a requirement for ideographic models. Prior studies systematically tested and validated theories of color compatibility and quantitative models for the aesthetic evaluation of color palettes in such design communities (O'Donovan, 2015; O'Donovan et al., 2011). Figure 1 shows top palettes from these online communities, with millions of user-created entries.

---

[1] https://color.adobe.com/explore

[2] https://colorlovers.com/palettes/

[Insert Figure 1 here]

The design task used in our model consists of individual agents who repeatedly attempt to create color palettes defined as the ordered arrangement of five adjacent color squares. Principles of color combination are used to structure this task based on a "color wheel", i.e., a circular ordering of color values organized by hue and brightness (Parkhurst & Feller, 1982). By means of linear regression, the most relevant features for predicting aesthetic evaluations of color palettes were previously identified as: hue probability, hue templates, hue entropy, and palette layout (O'Donovan, 2015). Thus, in this model we incorporate hue templates, hue entropy, and palette layout to inform the evaluation and synthesis of palettes generated and evaluated by design agents.

Hue templates are defined as fixed sets of rotations around the color wheel (O'Donovan, 2015, p. 27). Aesthetic preferences for one or two hues (monochromatic and dyadic combinations) are higher than more complex hue templates in such systems. Hue entropy is defined as a measure of the simplicity of a theme based on the distribution of its hue values so that entropy is lowest when all values are identical and highest when they are uniformly spread about the circle (O'Donovan, 2015, p. 30). Aesthetic preference is higher for mid-range entropy values in this type of systems. Lastly, palette layout (color ordering) shows an effect in aesthetic ratings in these systems and can be a straightforward way to improve the score of a palette by permutation of its colors (O'Donovan et al., 2011, p. 8).

We argue that this empirically-based multi-criteria approach provides a valid source for our agent-based model to be used for the exploratory analysis (Davis & O'Mahony, 2019) of a creative system (Csikszentmihalyi, 2014). Other arguably important features of the system are modelled stochastically or are treated as externalities. In a personal interview with the founder of one of these sites, we learned that factors that shape the creation and evaluation of palettes are likely to include the name of a palette and external events (current affairs, news, celebrities), as well as the social reputation of the creators (Monsef, 2019).

The current model integrates the invariants of design-task environments (Goel & Pirolli, 1992) as follows: design agents have incomplete access to the evaluation criteria; some constraints are enforced such as mid-range entropy and simpler hue templates, but how agents apply these constraints is individually determined and influenced by what other agents produce; evaluation includes conflicting criteria and some small changes can yield large evaluation effects such as by changing the contrast criterion by modifying a single color in the palette; high scores can be achieved in many possible ways; feedback is indirect and delayed in that agents only know the score of their own palettes after they create them and only know the score of other palettes if and when they use them to change their designs. This provides thus a simple yet sufficiently rich model that allows us to formulate and closely examine the core question of this study: How can systems *like this* (Kitcher, 2000, p. 66) help examine the premises underlying Big-C/little-c thinking?

## The Color Palette agent model

The agent-based system is presented here in pseudocode and graphically illustrated -the source code is publicly available online and by request from the first author. At setup time, the model defines a population of N agents as arrays of six variables: An integer value (`iD`) for identification, an integer value (`skill`) to define the individual *skill* level of each agent, an array (`colors`) of five color values in HSB space representing the five-color palette created by the agent, a coordinate variable (`xy`) to position the agent on the screen, an integer value (`score`) of the score calculated for the agent's palette, and an integer value (`frust`) representing the *frustration* level of each agent.

Agent(int iD, int skill, colors(size 5), coord xy, int score, int frust)

A simulation can start by initializing palettes and agents' characteristics at random, or they can be hand-picked for example by having all agents start with the same palette or specific proportions of low and high-skilled agents. To account for the stochastic nature of these models, each experimental condition is run thousands of times to obtain representative outcomes.

The agent variable `skill` defines a level of proficiency in the design of color palettes and it can be randomly initialized individually or uniformly assigned to the population at setup time. The `skill` of an agent can remain constant, or it can vary in response to its behavior -i.e., *learning*. In the current version, agents are individually assigned a `skill` value from 1 to *n* where *n* is the number of colors in the palette. An agent's `skill` defines how many entries in `colors` it can modify in its palette in every turn (simulation step). The individual variable `frust` is initialized at setup time for all agents as zero, and gradually increased every time that the agent sees no improvement in its palette's `score` after modifying its colors. When `frust` reaches a threshold, an agent is allowed to replace one value of its palette `colors` to a random hue and saturation values thus enabling a type of *creative mutation* to avoid convergence and escape local maxima (Boden, 1994, p. 528; Hofstadter, 1995). `frust` can increase at an individualized or a group rate depending on the question driving the modelling process. In the current version, `frust` increases uniformly and when an agent reaches a `frustT` threshold constant, it replaces one of its `colors` with a random color and its `frust` resets to 0. Weights can be assigned to the evaluation criteria, which can vary over a simulation to reflect externalities such as fashion trends, or, as in this model, remain constant during a simulation.

There are three instances in the model where randomness is used to account for externalities: at setup time in the initialization of the agent population and their initial color palettes; at every step to pair two agents to interact with each other; and occasionally, when an agent reaches a `frustT` threshold which enables it to introduce a random `color` on its palette. At every step of the simulation, two agents are selected at random to interact: one is labelled `agA` and the other `agB`. `agA` chooses at random one of its palette's `colors` to manipulate at this turn. If `agB's` palette has a higher `score`, then `agA` modifies its own palette to try and create a better (higher scoring) design as described below. First, we define how palettes are evaluated.

The evaluation function of color palettes in this system applies multiple criteria to account for the invariants of design-task environments (Goel & Pirolli, 1992). In this version of the model, three criteria are implemented: hueTemplate, hueEntropy, and layout following (O'Donovan et al., 2011). Hue templates are defined as monochromatic, dyadic, or triadic. Here, the `hueTemplateScore` is defined as inversely proportional to the angle covered by at least 4/5 of the hue values in the color wheel (360°) with a maximum score of 180 points for hue ranges ≤15°. This criterion as shown in Figure 2 rewards designs that combine a range of hues in a narrow segment of the chromatic space, and penalizes palettes that are more 'disperse'. Namely,

Array hueValues = (5 ordered color hues)

Do this four times with arc inversely proportional to the scoring of palettes

(monoScore, dyadScore, and triadScore)

1st: arc = 90, monoScore = 45, dyadScore = 40, triadScore = 15

2nd: arc = 45, monoScore = 90, dyadScore = 80, triadScore = 20

3rd: arc = 30, monoScore = 130, dyadScore = 120, triadScore = 25

```
    4th: arc = 15, monoScore = 180, dyadScore = 170, triadScore = 35
if (hueValues[0] - hueValues[3] <= arc) || (hueValues[1] - hueValues[4] <= arc)
    hueTemplateScore = monoScore;
if (hueValues[0] - hueValues[2] <= arc) && (hueValues[3] - hueValues[4] <= arc)
    hueTemplateScore = dyadScore;
if (hueValues[0] - hueValues[1] <= arc) && (hueValues[2] - hueValues[4] <= arc)
    hueTemplateScore = dyadScore;
if (hueValues[0] - hueValues[1] <= arc) && (hueValues[3] - hueValues[4] <= arc)
    hueTemplateScore = triadScore;
```

[Insert Figure 2 here]

We implement hue entropy via `hueEntropyScore`, proportional to a measure of variance of the hue values in the color palette. This criterion rewards designs that include more than one hue pair in different regions of the chromatic space as shown in Figure 3. Namely,

```
For (float num : hueValues)
  stDev += Math.pow(num - mean, 2);
  hueEntropyScore = Math.sqrt(stDev/ng);
```

[Insert Figure 3 here]

Color layout (`layoutScore`) is calculated to maximize harmonious sequences of hue and saturation values in a palette (O'Donovan et al., 2011). It is directly proportional to the number of hue values that are arranged in either a decreasing or an increasing sequence, plus the number of saturation values that are arranged also in a decreasing or increasing sequence. This 'sorting' criterion rewards designs where the colors form an ordered sequence in the chromatic space, as shown in Figure 4. Namely,

```
layoutScore += (hue1 > ... > hue5 ) || (hue1 < ... < hue5 );
layoutScore += (sat1 > ... > sat5 ) || (sat1 < ... < sat5 );
```

[Insert Figure 4 here]

These three evaluation criteria reward different qualities of a color palette and can be in contradiction since some designs (but not all) that receive high `hueTemplateScore` can score low on `hueEntropyScore`, etc. This allows for palette designs to score high in a variety of ways such as by maximizing one criterion over the rest or by balancing multiple criteria. This evaluation setup will also enable in future versions of the model to account for individual preferences, group formation, and temporal trends at the population level.

When two agents (`agA-agB`) are paired to interact, `agA` starts by choosing at random one of its palette's `colors` to manipulate, defined as `colorPos`. The first step that `agA` takes is to inspect the color layout of other agents in the group and modify its own color layout if its `score` is lower, as this can be considered the easiest move to try and improve the palette's `score`. Future work in this project will include ideation sessions where a group of (human) participants are asked to create color palettes in order to base these operators in data. If `agB`'s palette has a higher layoutScore, then `agA` swaps two of its colors' positions by selecting a random color to swap with the color in `colorPos`. Figure 5 illustrates how an agent modifies its palette's layout seeking to increase `layoutScore` by swapping two of its colors. Namely,

```
if (agB.layoutScore >= agA.layoutScore )
    color swapA = agA.colorPos;
    color swapB = agA.colorPos + 1;
    agA.colorPos = swapB;
    agA.color(swapB) = swapA;
    scoreUpdate(palette);
```

[Insert Figure 5 here]

The second step that `agA` takes is to modify the hue value of `colorPos` if it detects that `agB`'s `hueTemplateScore` is higher. In such case, `agA` modifies its palette's hue values by marginally approaching the hue value in `agB`'s corresponding `colorPos`. The rate in which `agA` makes this change is defined by parameter `dh` which we keep here constant as a value of 1 to guarantee that agents traverse the design space in small increments. Figure 6 illustrates how agent `agA` modifies the hue value of `colorPos` approaching `agB`'s palette. Namely,

```
if (agB.hueTemplateScore >= agA.hueTemplateScore )
   if (hue(agB.color1) > hue(agA.color1)) agA.color1 = color( min(360,
hue(agA.color1) + dH), saturation(agB.color1), brightness(agB.color1));
      scoreUpdate(palette);
```

[Insert Figure 6 here]

The third and last step that `agA` takes in this model is to modify the hue value of `colorPos` if it detects that `agB`'s `hueEntropyScore` is higher. In such case, `agA` translates the hue value `agA` steps in `colorPos` in the direction of the hue value in a neighbor of `colorPos` of its own palette. Figure 7 illustrates this process. Namely,

```
if (agB.hueEntropyScore >= agA.hueEntropyScore )
   if (colorPosition == 0) (
      if (hue(agA.color2) > hue(agA.color1)) agA.color1 = color( min(360,
hue(agA.color1) + dH), saturation(agA.color1), brightness(agA.color1));
         scoreUpdate(palette);
```

[Insert Figure 7 here]

Table 1 summarizes all the model variables described earlier, for clarity:

| **Variable name** | **Variable Type** | **Description** |
|---|---|---|
| `ID` | Integer | Identifies every agent |
| `skill` | Integer | Defines the number of colors an agent can change in their palette |
| `colors` | Array of color values | The color palette designed by every agent |

| | | |
|---|---|---|
| `xy` | Two-dimensional coordinate | Defines agent’s position in a display grid |
| `score` | Integer | Score (fitness) of a color palette |
| `frust` | Integer | Probability of a random change (mutation) in a palette |
| `agA and agB` | Agents | Designer agents that make up a population |
| `hueTemplateScore` | Evaluation criterion | Criterion applied in the evaluation of palettes based on hue templates |
| `hueEntropyScore` | Evaluation criterion | Criterion applied in the evaluation of palettes based on hue entropy |
| `layoutScore` | Evaluation criterion | Criterion applied in the evaluation of palettes based on palette’s layout |
| `colorPos` | Index | The position of a color in a palette |
| `simLimit` | Index | Number of steps that simulations run |
| `skill_1` | Low-skill agent | Agent with lowest `skill` value of 1/5 |
| `skill_5` | High-skill agent | Agent with highest `skill` value of 5/5 |
| `Hybrid` | Agent population | Group with minority of `skill_5` agents and majority of `skill_1` agents |

## Baseline conditions

To establish a baseline for the model, it is run with all agents having low skills (condition `skill_1`), namely when all agents can only change a single color of the palette each time they work on their design, and with all agents having high skills (condition `skill_5`), namely when all agents can change all the colors of their palette in their turn. Every other parameter of the simulation is kept constant across cases including random number seeds, rates and threshold of frustration, and scoring and tools to modify palettes. The aim is to define the outcomes with these two “extreme” conditions to establish a baseline to compare a version of the model with only a small minority of “genius-level” agents and a majority of “garden-variety” agents, as Big-C/little-c thinking assumes.

A total of 20 cases are run, each with a unique random seed and a constant `simLimit`, the number of total interactions between agents before the simulation ends. These were established by running cases for up to 1,000,000 steps with N = 10, 30, and 100 to confirm the findings reported. The results discussed here are with a group of 30 agents, `simLimit` = 50,000. Groups with all `skill_1` agents create an average of 132 top palettes, with an average score across cases of 393.82 and average maximum score of 431.59. In `skill_5` groups, agents tend to create around 10% more top palettes (143) with nearly 70% increase in number of times that top scores are reached. Figure 8 shows average top scores for 20 cases with `skill_1` (solid orange) and average top scores for 20 cases with `skill_5` (solid blue), as well as the highest top score across cases for `skill_1` (dash orange) and `skill_5` (dash blue).

Figure 9 shows the number of "hits" for each agent in the group: the instances across 20 cases that agents create a color palette of high value. To define the set of "best designs" we discard all initial palettes generated at model setup and only count high-scoring palettes produced after timestep 1. In `skill_1` conditions, (orange bars) agents generate a mean of 4.4 hits -standard deviation 2.08, while `skill_5` groups (blue bars) increase to 4.8 hits per agent and a higher variance (2.3). As expected, groups with lower skills produce fewer palettes and of lower quality compared to groups with higher skills where more palettes of higher quality are generated by a more select subgroup of agents.

[Insert Figure 8 here]

[Insert Figure 9 here]

Next, the Big-C/little-c distinction is examined with this model by creating groups that include a few "genius-level" agents with very high skills (`skill_5`, a creative elite) and a majority with very low skills (`skill_1`, the rest of the agents). According to Big-C/little-c thinking, in a system *of this type* only exceptionally skilled agents would design exceptionally good color palettes, and the rest of them with lower capacities would only (or mostly) generate incremental designs. For the little-c and Big-

C distinction to exist, a qualitative difference should be observable between the outputs generated by agents based on their individual capacities.

## Results

These results report 20 cases in N=30 groups, `simLimit` = 50,000 where only five agents have high design skills `skill_5` (agents #13 to #17) while the rest have low skills `skill_1` (agents #1 to #12 and #18 to #30). In this `hybrid` condition, we use the tag Big-C to refer to the elite `skill_5` agents and little-c to refer to the majority of `skill_1` agents in the group. An average total of 155 top palettes are created by these groups, an increase from both all `skill_1` (132) and all `skill_5` (143) conditions. The average hits per agent increases to 5.5, with a clear advantage of 17.2 top palettes for Big-C agents, and a significantly lower 2.76 for little-c agents as shown in Figure 11. As expected, these results show a concentration of top results in a few hands (the creative elite seeded in these groups), but arguably less expected is that `hybrid` groups are more likely to produce a higher number of top palettes than groups with all `skill_5` agents. To reiterate, `hybrid` groups perform better than uniform groups, even when their aggregate skills are markedly lower -a result that suggests the advantage of skill diversity in teams and calls for a separate analysis in the future.

The elite Big-C agents in `hybrid` groups generate 14 out of the 16 "Historically Best Designs" (`HBD`, defined as entries that reach 480 points or more in these conditions). In the top quartile of the 155 best palettes, Big-C agents show a clear dominance with 32 of the 40 entries -the remaining eight agents (all little-c) labelled `OTW` for "one-time wonders". Of the 155 best designs, Big-C agents generate a total of 86 entries (55%), and four of the little-c agents fail to generate a single top entry across all 20 cases. Therefore, the results in this scenario could be considered to confirm the

Big-C/little-c distinction to some extent: Big-C agents are *nearly* always behind the highest-scoring entries, and more often than not behind a top design. But not always. A closer look offers details worth examining. As noted, two of the 16 `HBD` entries are created by little-c agents, and 8 designs in the top quartile are also created by little-c agents. In addition, a rather unexpectedly large 45% of all 155 best palettes across cases are the work of these little-c agents too, a proportion that decreases to 28% of above-average top palettes and increases to 61% of below-average top palettes. `Hybrid` groups seem to owe their efficiencies to the designs contributed by *both* the highly productive minority as well as by low-ability individuals.

[Insert Figure 10 here]

The considerable contributions of little-c agents to the top scores indicate that if the Big-C label is applied solely based on outcomes, then in many instances `skill_5` agents would not be recognized as Big-C agents, and many `skill_1` agents would be inferred to be Big-C agents. In other words, solely based on the scoring of palettes, it is unfeasible to distinguish between `skill_1` and `skill_5` agents. This emphatically shows the conceptual weakness of labelling a creator based on their output alone without directly accounting for generative processes. To reiterate, in systems of this type a number of agents with significantly lower skills are able to generate top designs including some reaching the top historical achievements.

These results indicate that the computational examination of the Big-C/little-c categories shows that it is a conceptually weak construct because it cannot consistently explain the characteristics of the generative processes of creators. In these models, only 55% of the time the performance of high achievers is explained by individual factors, while in 45% of cases there is no intrinsic individual causation. This study thus confirms *how* Big-C can be a quality that is attributed by observers of the outcomes which hides key attributes of how individuals create. The main conclusion from this work is that higher capacities are neither necessary, much less sufficient, to explain exceptional outcomes in systems *of this type* (Kitcher, 2000, p. 66).

## Discussion

This study set to conduct a formal examination of a distinction about levels of creativity that has been highly cited despite its vague rhetorical framing. An in-depth inquiry of the Big-C and little-c hierarchy of creativity is conducted via a nomothetic agent-based simulation that forces researchers to be explicit about the premises, variables, and parameters in ways that enable building theory that is "decidedly falsifiable" (Jackson et al., 2017). Simulation has been defined as a third way of building theory: starting with a deductive approach to grow data that is analyzed inductively (Axelrod, 2013) to generate or grow the phenomena of interest (Epstein, 1999). The patterns emerge from clearly specified rules rather than from real-world evidence making simulation valuable to aid intuition and perform computational thought experiments (Axelrod, 2013). The validity of this model to support exploratory analysis about the Big-C/little-c distinction depends therefore on several factors: first, the extent to which it reflects the key qualities of the target system including the behavior of the agents, the type of outputs designed, and the mechanisms used to evaluate and synthesize designs. Second, the appropriate use of pseudo-random parameters to account for externalities. Thirdly, the causation mechanisms that give rise to the different types of outcomes given the parameters in the baseline and the experimental conditions. Lastly, the adherence between the description of the model and its implementation in computer code, which is shared publicly for other researchers to verify and recreate the data presented in this paper.

Whilst the agent model does include stochastic decision-making, we argue that the level of randomness employed accounts for the type of externalities that can shape the behavior of such systems. To the extent that this agent model captures some of the processes that characterize creative behavior in fields such as the design of color palettes, these results can help to better understand how the Big-C/little-c distinction is a "false dichotomy" (Runco, 2014). By showing that not all achievements are created by agents who possess advanced skills or superior traits, these results demonstrate the conceptual weakness of dividing creators in categories such as "mavericks" and "everyday people" (Kaufman & Beghetto, 2009). These labels have their origins in societal dynamics and therefore it can be problematic to use them to characterize inherent qualities of individual processes or people.

Hierarchies of creativity have been argued rhetorically based on whether creations are evaluated by others as amateur, accomplished-but-not-yet-eminent, or genius-level (Kaufman & Beghetto, 2009; Merrotsy, 2013). Our study shows that in social systems like the creative communities modelled here, although high-skilled individuals are likely to concentrate high-quality creations, the boundary between high and low achievers is not as clear as the Big-C/little-c distinction would require. Some agents with high skills do not perform as well as others, and some agents with much lower skills can match and surpass the outcomes of agents with clear advantages. At the heart of this work lies the concept of "creative situations" (Gero & Sosa, 2002). Namely, some situations can favor certain creators over others regardless of their individual merits, such as accessing information or their creations being recognized first, or pure chance (Simonton, 2004). In complex social systems, causation transcends the individual sphere and includes group and social factors over time, thus supporting the call for creativity to expand individualistic approaches to address situational causation (Mumford, 1995; Runco, 2008).

In systems like this, agents can build upon the solutions of others. A similar mechanism of co-creation has been documented in the collaboration of designers in creative teams where some specialize in giving ideas and others in taking and improving on them, without implying that one type is more creative than the other (Elsbach & Flynn, 2013). Thus, the underlying principles captured in this study may be applicable to real-world design systems inasmuch as designers form their ideas informed by the ideas of others. More work needs to be done to test this in more veridical agent-based models of creativity, including those where the design task captures the dynamics of a complex real-world design domain.

Further work is also needed to include learning mechanisms for agents to account for the evidence from studies of creativity and expertise. For example, three distinct agent mechanisms can be defined in future modelling to represent divergent thinking, intelligence, and expertise and their effects on ideation (Vincent et al., 2002). The need for conceptual clarity in how such constructs are implemented in artificial creators will require a close examination of how they have been defined and assessed in such studies.

Even in conditions where individuals may seem to create in isolation, when designs become available in the marketplace, they can and do influence other creators. Thus, we argue that the modelling of creativity needs to account for the link between individual and group-level factors since for any act of creation, individuals rely on the ideas and means produced by others including their peers, competitors, collaborators, and "the shoulders of giants" who came before (Sosa & Gero, 2016). A flaw behind the Big-C/little-c distinction is that it conflates levels of creativity by attributing a label based on evaluation to stereotypes of individual creators (Mumford et al., 2008; Sosa & Gero, 2016).

In design studies it has been speculated that "it is entirely possible that creativity is rather common in everyday design projects and is not just the province of the activity of exceptional designers" (Wiltschnig et al., 2013). The study presented here supports this way of seeing creativity, and further provides support to avoid classifying designers into "everyday" and "exceptional" categories in the first place.

The refutation of the Big-C/little-c distinction presented here could easily be misconstrued as an argument for the fallacy that "everyone can be Mozart". Universal and plural views of creative capacities (Arendt, 2013) allow for a wide range of abilities including domain-specific talents, and a wide range of environmental advantages. Nature and nurture interplay to make certain creators better equipped, privileged, and positioned to generate exceptional outcomes in all sorts of fields, professions, and occupations. Rather, the refutation of Big-C/little-c creativity offered in this paper shows that explanations of exceptional creators require both individual and situational factors. An infrastructure was in place for Mozart's talents to be informed, developed, applied, and valued. Musical genres and notation, instruments and performers, patrons, a public, all of these and other essential elements need to be accounted for when referring to the "Mozart" appellative, and not only the individual creator.

The evaluation of a creative outcome cannot be reduced to the process and person who produced it, as Big-C/little-c thinking requires. Not only it is conceptually incorrect to characterize a generative process or a person based on a social evaluation of the outcome, it is also contrary to the principle of parsimony and is based on circular reasoning that assumes that generative processes of creation are of a different type (Dietrich, 2007; Weisberg, 1993). An entailment of this refutation is that acclaimed designs need not be created equally to other acclaimed designs, and they need not be created differently than unacclaimed designs. This allows for scenarios where designs of equal or higher value go unnoticed for long periods, some of them forever lost in history or re-discovered at a later stage (Simonton, 2004).

Future work with agent-based simulations of creativity will address the increased productivity of hybrid groups reported in our findings. We are also working on adding roles and mechanisms to these models to address other research questions for the study of creativity including the role of influential designs and the changing nature of *taste* which can affect the evaluation criteria of color palettes in this model. The source computer code is publicly available online and upon request.

## References


Arendt, H. (2013). *The Human Condition*. University of Chicago Press.

Axelrod, R. (2013). Advancing the art of simulation in the social sciences. In R. Conte, R. Hegselmann, & P. Terna (Eds.), *Simulating Social Phenomena* (pp. 21-40). Springer Science & Business Media.

Boden, M. A. (1994). Précis of the creative mind: Myths and mechanisms. *Behavioral and Brain Sciences, 17*(3), 519-531.

Boghossian, P., & Peacocke, C. (2000). *New essays on the A Priori*. Clarendon Press.

Calder, M., Craig, C., Culley, D., de Cani, R., Donnelly, C. A., Douglas, R., Edmonds, B., Gascoigne, J., Gilbert, N., Hargrove, C., Hinds, D., Lane, D. C., Mitchell, D., Pavey, G., Robertson, D., Rosewell, B., Sherwin, S., Walport, M., & Wilson, A. (2018). Computational modelling for decision-making: where, why, what, who and how. *Royal Society Open Science*, *5*, 172096. https://doi.org/10.1098/rsos.172096

Carley, K. M. (2019). Social-Behavioral Simulation: Key Challenges. In P. K. Davis, A. O'Mahony, & J. Pfautz (Eds.), *Social-Behavioral Modeling for Complex Systems* (pp. 741-752). John Wiley & Sons Inc. https://doi.org/10.1002/9781119485001.ch32

Csikszentmihalyi, M. (2014). *The systems model of creativity : the collected works of Mihaly Csikszentmihalyi*. Springer.

Davis, P. K., & O'Mahony, A. (2019). Improving Social-Behavioral Modeling. In P. K. Davis, A. O'Mahony, & J. Pfautz (Eds.), *Social-Behavioral Modeling for Complex Systems* (pp. 15-48). John Wiley & Sons.

Dennett, D. C. (2013). *Intuition pumps and other tools for thinking*. WW Norton & Company.

Dietrich, A. (2007, 2007/05/01/). Who's afraid of a cognitive neuroscience of creativity? *Methods*, *42*(1), 22-27. https://doi.org/https://doi.org/10.1016/j.ymeth.2006.12.009

Elsbach, K. D., & Flynn, F. J. (2013, 2013/06/01). Creative Collaboration and the Self-Concept: A Study of Toy Designers. *Journal of Management Studies*, *50*(4), 515-544. https://doi.org/10.1111/joms.12024

Epstein, J. (1999). Agent-Based Computational Models and Generative Social Science. *Complexity*, *4*(5), 41-60.

Gero, J. S., & Sosa, R. (2002). *Creative design situations: Artificial creativity in communities of design agents* Computer-Aided Architecture Design Research in Asia CAADRIA, Malaysia.

Gilbert, N., & Ahrweiler, P. (2006). The epistemologies of social simulation research. International Workshop on Epistemological Aspects of Computer Simulation in the Social Sciences,

Goel, V., & Pirolli, P. (1992, 1992/07/01/). The structure of design problem spaces. *Cognitive Science*, *16*(3), 395-429. https://doi.org/http://dx.doi.org/10.1016/0364-0213(92)90038-V

Hofstadter, D. R. (1995). *Fluid concepts & creative analogies : computer models of the fundamental mechanisms of thought*. Basic Books.

Jackson, J. C., Rand, D., Lewis, K., Norton, M. I., & Gray, K. (2017). Agent-Based Modeling:A Guide for Social Psychologists. *Social Psychological and Personality Science*, *8*(4), 387-395. https://doi.org/10.1177/1948550617691100

Kaufman, J. C., & Beghetto, R. A. (2009). Beyond Big and Little: The Four C Model of Creativity. *Review of General Psychology*, *13*(1), 1-12. https://doi.org/10.1037/a0013688

Kitcher, P. (2000). A priori knowledge revisited. In P. Boghossian & C. Peacocke (Eds.), *New Essays on the A Priori* (pp. 65-91). Clarendon Press.

Merrotsy, P. (2013, 2013/10/01). A Note on Big-C Creativity and Little-c Creativity. *Creativity Research Journal*, *25*(4), 474-476. https://doi.org/10.1080/10400419.2013.843921

[Record #4508 is using a reference type undefined in this output style.]

Mumford, M. D. (1995, 1995/10/01). Situational Influences on Creative Achievement: Attributions or Interactions? *Creativity Research Journal, 8*(4), 405-412. https://doi.org/10.1207/s15326934crj0804_7

Mumford, M. D., Hunter, S. T., & Bedell-Avers, K. E. (2008). *Multi-level issues in creativity and innovation* (Vol. 7). Emerald Group Publishing Limited.

O'Donovan, P. (2015). *Learning Design: Aesthetic Models for Color, Layout, and Typography* University of Toronto]. Toronto.

O'Donovan, P., Agarwala, A., & Hertzmann, A. (2011). Color compatibility from large datasets. *ACM Trans. Graph., 30*(4), 1-12. https://doi.org/10.1145/2010324.1964958

Parkhurst, C., & Feller, R. L. (1982). Who invented the color wheel? *Color Research & Application, 7*(3), 217-230. https://doi.org/10.1002/col.5080070302

Rangoni, R. (2014). From Informal Thought Experiments to Agent-Based Models A Progressive Account of Modeling in the Social Sciences. In L. Magnani (Ed.), *Model-Based Reasoning in Science and Technology* (pp. 471-478). Springer.

Runco, M. A. (2008). Creativity research should be a social science. In M. D. Mumford, S. T. Hunter, & K. E. Bedell-Avers (Eds.), *Multi-Level Issues in Creativity and Innovation (Research in Multi-Level Issues, Vol. 7)* (pp. 75-94). Emerald Group Publishing Limited. https://doi.org/https://doi.org/10.1016/S1475-9144(07)00003-3

Runco, M. A. (2014). "Big C, Little c" Creativity as a False Dichotomy: Reality is not Categorical. *Creativity Research Journal*, *26*(1), 131-132. https://doi.org/10.1080/10400419.2014.873676

Runco, M. A., & Jaeger, G. J. (2012, 2012/01/01). The Standard Definition of Creativity. *Creativity Research Journal*, *24*(1), 92-96. https://doi.org/10.1080/10400419.2012.650092

Salvatore, S., & Valsiner, J. (2010). Between the general and the unique: Overcoming the nomothetic versus idiographic opposition. *Theory & Psychology, 20*(6), 817-833.

Simonton, D. K. (2004). *Creativity in Science: Chance, Logic, Genius, and Zeitgeist*. Cambridge University Press.

Simonton, D. K. (2017). Big-C Versus Little-c Creativity: Definitions, Implications, and Inherent Educational Contradictions. In R. A. Beghetto & B. Sriraman (Eds.), *Creative Contradictions in Education: Cross Disciplinary Paradoxes and Perspectives* (pp. 3-19). Springer International Publishing. https://doi.org/10.1007/978-3-319-21924-0_1

Sosa, R., & Gero, J. S. (2016, 2016/01/02). Multi-dimensional creativity: a computational perspective. *International Journal of Design Creativity and Innovation, 4*(1), 26-50. https://doi.org/10.1080/21650349.2015.1026941

Sternberg, R. J. (2018). What's Wrong with Creativity Testing? *The Journal of Creative Behavior, 53*(1), 121. https://doi.org/10.1002/jocb.237

Sun, Z., Lorscheid, I., Millington, J. D., Lauf, S., Magliocca, N. R., Groeneveld, J., Balbi, S., Nolzen, H., Müller, B., & Schulze, J. (2016). Simple or complicated agent-based models? A complicated issue. *Environmental Modelling & Software, 86*, 56-67.

Vincent, A. S., Decker, B. P., & Mumford, M. D. (2002, 2002/04/01). Divergent Thinking, Intelligence, and Expertise: A Test of Alternative Models. *Creativity Research Journal, 14*(2), 163-178. https://doi.org/10.1207/S15326934CRJ1402_4

Weisberg, R. W. (1993). *Creativity "Beyond the myth of genius"*. W H Freeman & Co.

Wiltschnig, S., Christensen, B. T., & Ball, L. J. (2013, 2013/09/01/). Collaborative problem–solution co-evolution in creative design. *Design Studies, 34*(5), 515-542. https://doi.org/https://doi.org/10.1016/j.destud.2013.01.002


Figure 1. Most *loved* entries in ColorLovers and most *popular* entries in Adobe Color as of July 2019

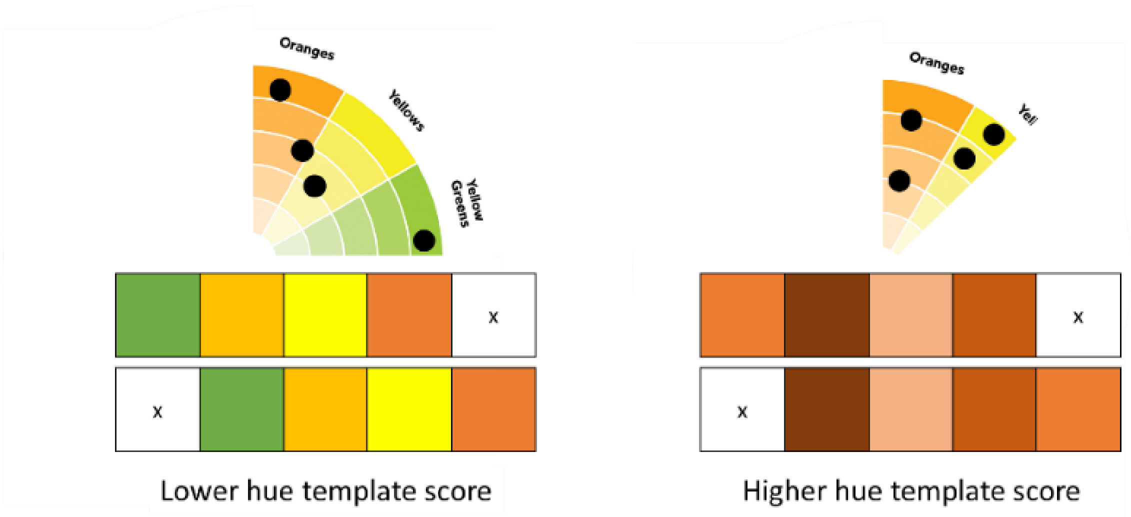


Figure 2. Sample hue templates that score lower (left) and higher (right) as most of its hues range from wider (left) or narrower (right) angles of the color wheel (Parkhurst & Feller, 1982).

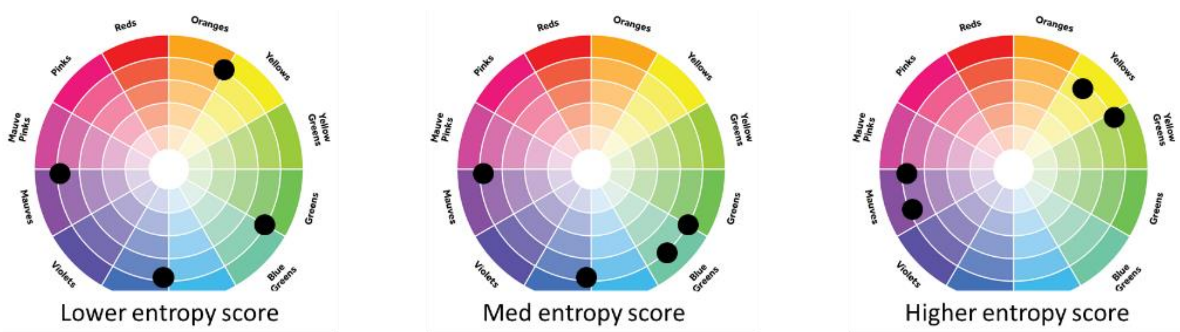

Figure 3. sample hue entropy scores from lower (left) to medium (centre) to higher (right) showing how the pairing of hue values improves the hueEntropyScore when hue pairs are spread across the color wheel.

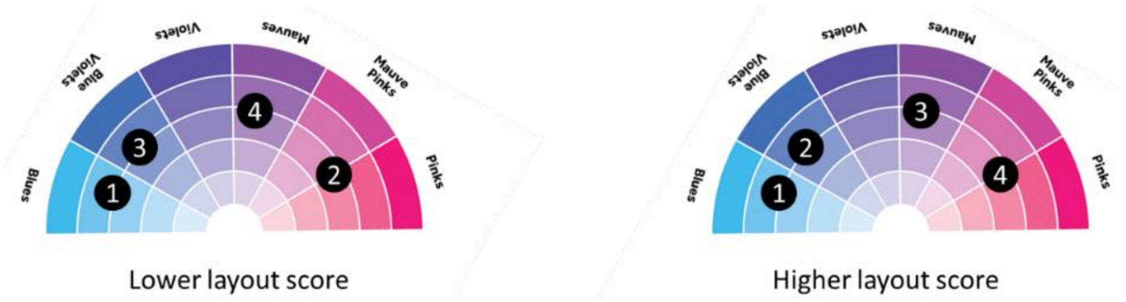


Figure 4. Sample layout scores from lower (left) to higher (right) as defined by the ordering of the palette colors, hue and saturation values are the same but positioned in a different ordered sequence.

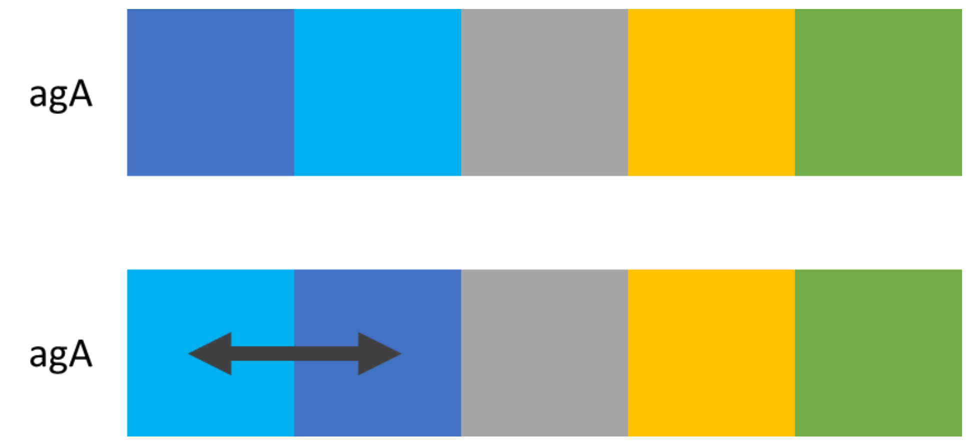


Figure 5. Agent agA swaps two of its colors seeking to increase layoutScore

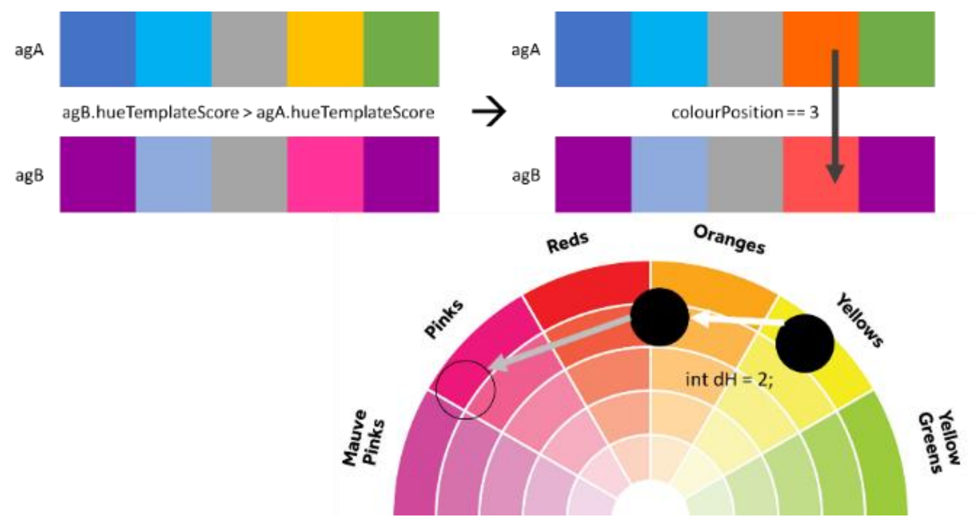

Figure 6. Agent agA modifies the hue of one of its colors seeking to increase hueTemplateScore

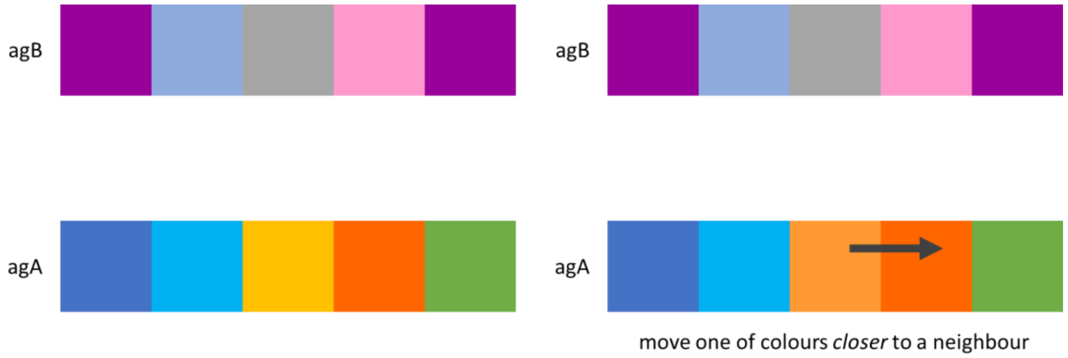


Figure 7. Agent agA modifies the hue value of colorPos to approach the hue value in a neighbour of colorPos of its own palette.

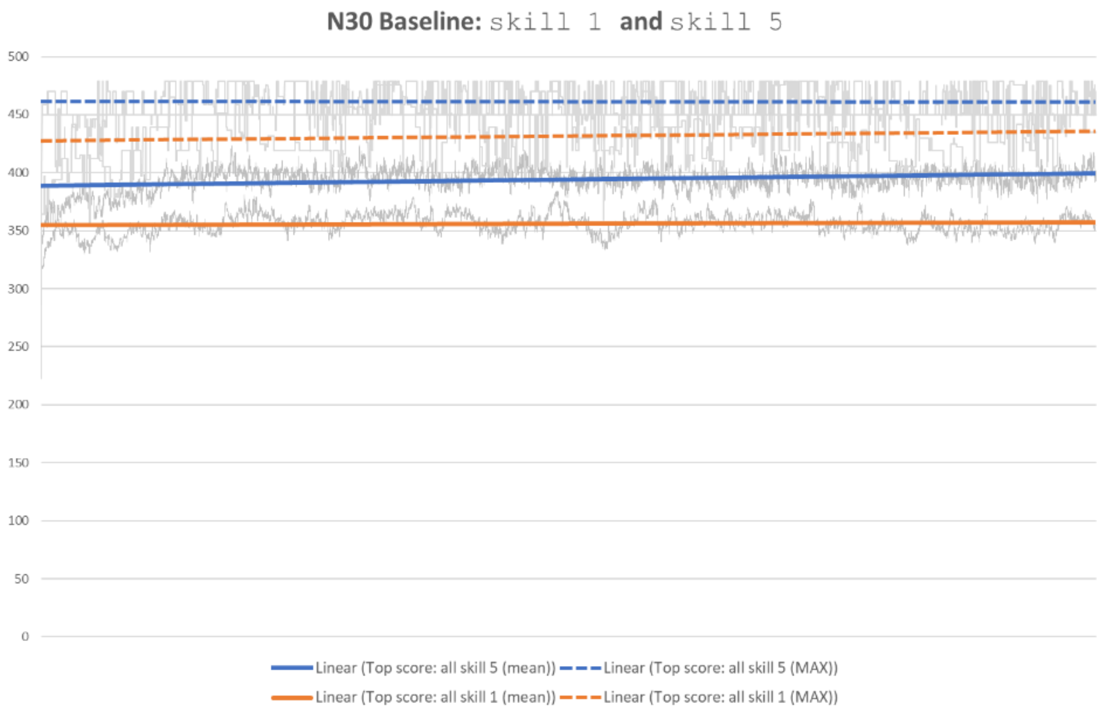


Figure 8. Mean scores and maximum scores in skill_1 and skill_5

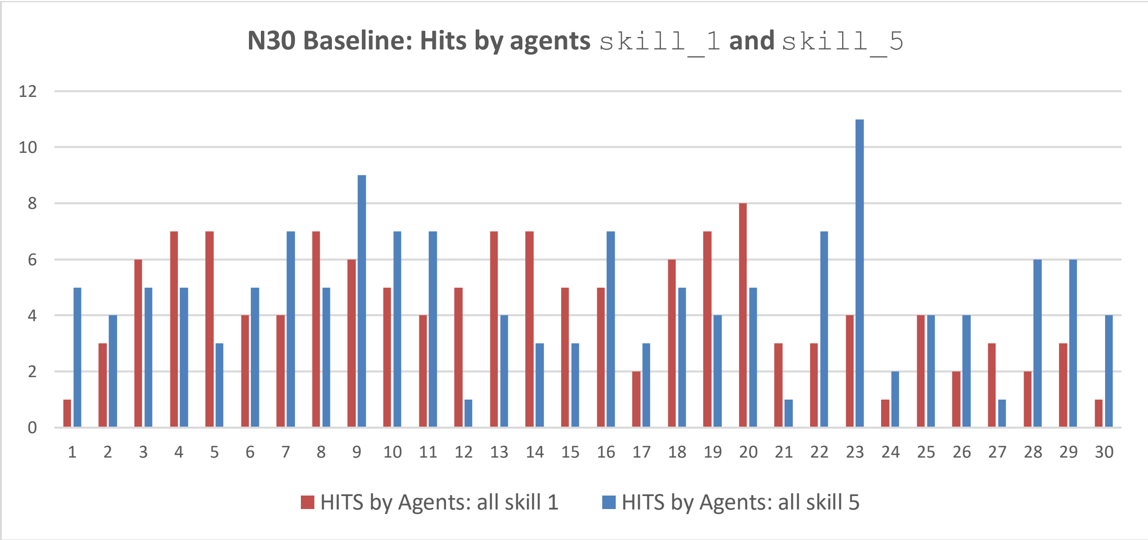


Figure 9. Hits by agents in skill_1 and skill_5

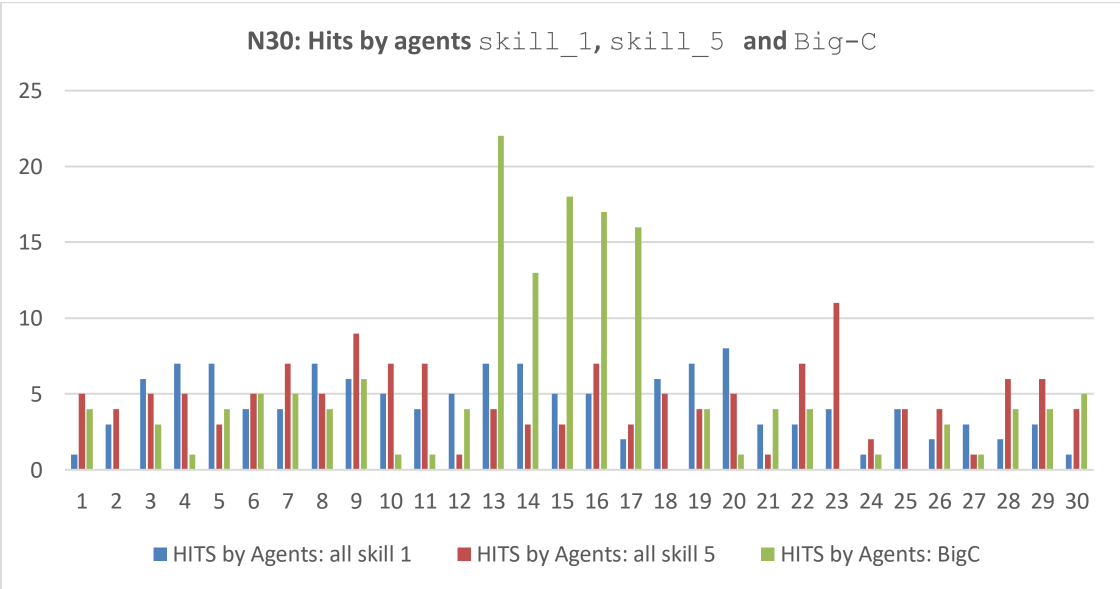


Figure 10. Hits by agents in skill_1 (blue bars), skill_5 (orange bars), and in scenarios with a minority of Big-C agents (#13 to #17) and a majority of little-c agents (the rest)