

Developing Creativity

Artificial Barriers in Artificial Intelligence*

Kyle E. Jennings**

University of California, Berkeley
Institute of Personality and Social Research
jennings@berkeley.edu

Abstract. The greatest rhetorical challenge to developers of creative artificial intelligence systems is convincingly arguing that their software is more than just an extension of their own creativity. This paper suggests that the key feature this requires is “creative autonomy,” which exists when a system not only evaluates creations on its own, but also changes its standards without explicit direction. Paradoxically, developing creative autonomy is argued to require that the system be intimately embedded in a broader society of other creators and critics. A sketch is provided of a system that might be able to achieve creative autonomy, provided it initially liked others in the society to different extents, and had to remain “proud” of its past work. This should lead to emergent dynamics that enable the creative AI to be more than a simple blend of the influences it was predicated upon, though this awaits empirical demonstration.

Key words: computational creativity, autonomy, socially-inspired computing

1 The Quest for Creative Autonomy

Much of the theoretical work in creative artificial intelligence tries to specify when a system has gone beyond simply doing the bidding of its programmer. For instance, one rationale for Boden’s [1] “transformational” criterion is that since the programmer creates the initial search space with a particular view of what is possible, a system that transformed that space would be going beyond the programmer’s vision. Ritchie’s [2] “inspiring set” helps determine whether an idea produced by the system was directly involved in the system’s creation or training. Finally, Colton’s [3] inclusion of imagination in his creative tripod

* To appear in *Proceedings of the Fifth International Joint Workshop on Computational Creativity*, Madrid, Spain, September 2008.

** The author is grateful to three anonymous reviewers, whose feedback greatly improved this work, and to the organizers of and participants in the 2008 AAAI Spring Symposium on Creative Intelligent Systems, in discussion with whom these ideas were hatched.

hearkens to the autotelic exploration of ideas that is not tethered to outside forces.

The difference between greater and lesser creativity lies not in how you solve problems, but rather in what problems you choose to solve [4]. Therefore, creative systems need to be seen as pursuing an independently chosen vision. This means that in addition to being able to independently apply the standards it knows, a system must be able to independently change the standards it uses. This ideal will be called “creative autonomy,” and represents the system’s freedom to pursue a course independent of its programmer’s or operator’s intentions.

Because some skeptic will always accuse creative AI of being no more than an elaborate tool for expressing the programmer’s creativity, any contact between the system and the programmer comes under suspicion. This can lead to the desire for systems that are hermetically sealed from the outside world. However, human creativity, which is clearly autonomous, takes place within a rich web of social interactions. It is in making sense of and responding to these interactions that we arrive at a style that is unique to us, yet still acceptable enough for others to take seriously. Creative autonomy will likewise be argued to emerge out of the interactions with multiple critics and creators, not from solitary confinement.

Criteria A system will be said to have creative autonomy if it meets the following three criteria:

Autonomous Evaluation the system can evaluate its liking of a creation without seeking opinions from an outside source

Autonomous Change the system initiates and guides changes to its standards without being explicitly directed when and how to do so

Non-Randomness the system’s evaluations and standard changes are not purely random

Autonomous evaluation requires that from the moment the system begins generating an opinion until the moment it has finished issuing the opinion, it does not consult another human or machine intelligence. However, the system is free to ask for opinions at other times, and to store this information. Autonomous evaluation could be easily achieved by using preprogrammed standards or by learning another source’s standards, both of which could be used to bootstrap a system. After this, however, autonomous change requires that the system be able to independently change its standards. Though external events may prompt and guide changes, the system cannot rely on another source to tell it when to change standards, or when its new standards are acceptable, nor can it simply make some fixed transformation to another source’s standards.

An easy way to satisfy both criteria would be to issue random decisions, or to make random changes at random times. The final criterion is meant to prevent this. Many algorithms incorporate randomness, so not all randomness is precluded. For instance, the system could resolve conflicts between standards randomly, or it could test random perturbations to its standards. Aside from special cases like these, however, it cannot simply make random decisions. Of course, this does not guarantee predictable outcomes.

Creative Autonomy and Creativity Being considered creative depends not just on what you create, but also on the circumstances in which you create [5]. “Creative autonomy” is meant to capture some of the circumstances that creative AI are seen to lack. A system with creative autonomy has the potential to produce creative results, in which case it could be called “creative”. However, as in our own lives, creative potential is not a guarantee of creative success.

The next section analyzes the information a system could use to issue autonomous evaluations, and looks at how that information changes through social interactions. The third section discusses some techniques that could be used to produce autonomous change in the context of these interactions. The final section summarizes this work and suggests implications and future directions.

2 Learning to Evaluate

This section describes how a creator could learn evaluation standards via its interactions with others. Since creators and non-creators both have opinions, these others will be called “critics” rather than “creators”. Though the notation used is based on Wiggins’ model of creative search [6], the only thing of interest is how a creator learns to judge the quality of finished products. How these standards affect the search process, as well as how they apply to intermediate products, are left as interesting questions for future work. This model assumes that there is at least one critic, though its more interesting features do not apply unless there are more. Though the processes described here are inspired by human creativity, they could be implemented in a society of solely AI creators, or in a mixed human-machine society.

Subjectivity Following [7], we will assume that a creation’s value is socially constructed, and that different critics have different standards. Wiggins uses \mathcal{E} for the knowledge representing these standards, which can be subscripted to indicate whose standards are in question, e.g., \mathcal{E}_i . Unlike in Wiggins’ model, no representation for this knowledge will be assumed.

In addition to knowing that people have different opinions, we can often estimate a typical or specific person’s opinion. Thus, the knowledge in \mathcal{E}_i can be segmented by whose evaluations the knowledge is about. For this we will use the subscript ij , where i is the perceiver and j is whose opinion is perceived. A dot (“.”) will represent the typical critic. Thus,

$$\mathcal{E}_i = \langle \mathcal{E}_{i.}, \mathcal{E}_{i1}, \dots, \mathcal{E}_{ii}, \dots, \mathcal{E}_{iN} \rangle$$

where N is the number of critics in the society. Knowing other critics’ preferences lets a creator target an audience, and so it is important for the information to be correct. Therefore, assume that creators continuously correct inaccuracies.

For sake of argument, assume that creators represent knowledge at the most general level possible and avoid duplicating knowledge. This means that something that applies to most creators would be stored in $\mathcal{E}_{i.}$, and creator j ’s deviations from this would be stored in \mathcal{E}_{ij} . If creator i knows nothing specific about

creator j 's standards, then $\mathcal{E}_{ij} = \emptyset$. We will assume the most difficult case in which creators start with no standards of their own, i.e., $\mathcal{E}_{ii} = \emptyset$.

Making Evaluations Creator i 's evaluation of creation c from critic j 's perspective is denoted by $E_{ij}(c)$. Though the notation is analogous, E_{ij} is not simply the application of \mathcal{E}_{ij} . This is because the applicability of \mathcal{E}_{ij} depends on c . For instance, if the discipline is furniture design, creator i might know a great deal about how j evaluates chairs, but nothing about how j evaluates tables. If c is a table, it would make more sense for i to rely on \mathcal{E}_i than \mathcal{E}_{ij} .

Wiggins writes $\llbracket \mathcal{X} \rrbracket$ to mean the translation of the knowledge in \mathcal{X} to a function from creations to real numbers in $[0, 1]$. We will extend this to map to $[0, 1] \times [0, 1]$, for the result and the confidence in that result. Additionally, we need a function that can aggregate different evaluations and confidence levels into a single answer. Heuristics such as assuming that people from similar backgrounds have similar opinions could compensate for missing information. These details don't matter here, and so we'll simply say that the system has background social knowledge, \mathcal{S}_i , and a function F_i that uses it to consolidate the other information. Given this, for $i \neq j$ we have:

$$E_{ij}(c) = F_i(\mathcal{S}_i, j, \llbracket \mathcal{E}_i \rrbracket(c), \llbracket \mathcal{E}_{i1} \rrbracket(c), \dots, \llbracket \mathcal{E}_{iN} \rrbracket(c))$$

In the extreme case, a creator's own opinion would only depend on knowledge in \mathcal{E}_{ii} . However, by hypothesis \mathcal{E}_{ii} is initially empty, meaning that the creator must construct its opinion from what it knows about others' opinions. Though we could make the system issue the most representative opinion, it will prove more interesting if it prefers to emulate some critics more than others. These affinity levels are stored in \mathcal{A}_i , and are discussed in the next section. We can now define an analogous function to F_i :

$$E_{ii}(c) = F'_i(\mathcal{S}_i, \mathcal{A}_i, \llbracket \mathcal{E}_i \rrbracket(c), \llbracket \mathcal{E}_{i1} \rrbracket(c), \dots, \llbracket \mathcal{E}_{iN} \rrbracket(c))$$

Note that E_{ii} would be just one component of the creator's objective function during search (cf. [8]), but is the only function creator i uses to evaluate its own and others' finished products.

Communication Creator i learns to make autonomous evaluations via interactions with other critics. Suppose that a creator i has made a creation c , which is observed by a critic $j \neq i$. There are three broad classes of information that can be communicated.

Evaluation A simple "like/dislike" judgment. Creator j communicates $E_{jj}(c)$, and then creator i adjusts its knowledge until $E_{ij}(c) \approx E_{jj}(c)$.

Correction Critic j creates c' , a modification of c that it likes better. Creator i updates its knowledge so that $E_{ij}(c') > E_{ij}(c)$, and tries to determine what changes between c and c' increased j 's liking.

Criticism Justifications for an evaluation or correction, e.g., what is pleasing or what criteria were used. Creator j communicates knowledge in or derived from \mathcal{E}_j to creator i , which attempts to integrate this into \mathcal{E}_i . If i cannot make $E_{ij}(c) \approx E_{jj}(c)$, then i might ask j for clarification.

In each case, creator i adjusts \mathcal{E}_i in order to reproduce j 's evaluation. Because knowledge is represented at the most general level and duplication is avoided, this should always result in change to $\mathcal{E}_{i.}$ or to \mathcal{E}_{ij} . These processes cannot by themselves make \mathcal{E}_{ii} non-empty.

In creative AI systems that only allow interaction with one critic (the programmer), all of the system's knowledge can be represented in $\mathcal{E}_{i.}$, meaning that $E_{ii}(c) = E_{ij}(c) = E_{i.}(c) = \llbracket \mathcal{E}_{i.} \rrbracket(c)$, i.e., the system parrots back its understanding of the critic's standards. The situation improves somewhat with multiple critics since the system forms E_{ii} from many different sets of standards in ways dependent on \mathcal{S}_i and \mathcal{A}_i . However, it still only offers direct translations of other critics' standards. What's more, in both cases, the system only changes its standards in reaction to and in proportion to changes in other critics' standards. Hence, though these processes support autonomous evaluation and are non-random, they are not enough for creative autonomy. The next section suggests some extensions that would add the missing component, autonomous and non-random change.

3 Changing Standards

If the system faithfully updates its knowledge of others' standards, autonomous change will not occur until there is knowledge in \mathcal{E}_{ii} . Since all of the system's knowledge comes from other critics and is stored at the most general level, there is as yet no reason for this to happen. Inspired by human psychological processes that would be simple to implement, this section suggests some reasons that \mathcal{E}_{ii} might be initially populated and subsequently changed.

3.1 Additional Behaviors

As described so far, the system combines others' preferences according to how applicable they are and how much it "likes" each critic. This section first describes how "liking" could be initially configured and then changed. Thus far the system never has cause to doubt its own evaluations. One such reason will be introduced, which will later be argued to lead to including knowledge in \mathcal{E}_{ii} .

Affinity Any number of rules could be used to set the initial affinities in \mathcal{A}_i , all of which have a basis in human psychology:

Propinquity Our friendships [9] and collaborations [10] are largely determined by physical proximity. Analogously, the system could be initially set to prefer creators who are nearby in some topology, real or imposed.

Similarity We subconsciously favor people with similar backgrounds. In a society of artificial creators with varied parameterizations, similarly parameterized creators might initially prefer each other.

Popularity When we cannot make sense of a speaker’s message, we decide whether to believe her based on cues about her prestige [11], e.g., age (time in the society) or popularity (received affinity).

Some affinity changes would be independent of the system’s evaluations:

Familiarity Absent other discernable differences, we tend to prefer people and things we have seen before [12]. Frequent interactions could increase liking.

Mutual Affinity We are more apt to like someone if they first show that they like us [13]. The system could increase its affinity for critics that evaluate the system’s creations positively.

Finally, affinity could adjust in response to the creator evaluating a critic’s work, or by how closely the creator and critic agree on evaluations of a third creator’s work. At first this would not be very meaningful, but as the creator absorbs influences and gains independence it should lead to less tautological changes.

Pride Unsure about the quality of their work, novices are particularly sensitive to praise and criticism. The sting of failure can be offset by the memory of success, making preserving good memories important.

This could be modeled by storing a memory of past successes, \mathcal{M}_i , and their last average evaluation, $M_i = \sum_{c \in \mathcal{M}_i} E_{ii}(c) / |\mathcal{M}_i|$. Only highly salient creations would be stored (ones that elicited “pride”), such as ones that got unexpectedly high evaluations (relative to recent creations, other creators’ creations, or the critic’s typical evaluation), particularly from a critic the creator likes. As with a person who concludes that all of her prior work was worthless, there could be negative repercussions for the system if the value of M_i suddenly dropped. As discussed next, avoiding this could lead the system to develop its own standards.

3.2 Bootstrapping and Changing \mathcal{E}_{ii}

This section introduces three processes that could introduce and change knowledge in \mathcal{E}_{ii} . As before, each is inspired by human behavior. They are sketched here, and discussed relative to creative autonomy in the next section.

Cognitive Dissonance Consider a novice whose evaluations mirror an influential mentor’s, and whose self-confidence rests on the memory of past successes. Suppose that one particular creation, which was highly rated by his mentor, is a large source of pride. He only understands why that work was good in terms of how he understands his mentor’s preferences, which he trusts since he respects that mentor. Now suppose that the mentor strongly criticized a highly similar creation, throwing into doubt his understanding of the mentor’s standards. Or, perhaps an unrelated event would make him lose respect for the mentor, leading him to discount the mentor’s opinion. In either case, he could no longer justify such a high evaluation for his prized creation, leading to a dilemma: believe the reduced evaluation, and hence that he’s not as good as he thought; or, doubt the

new evaluation, and continue to believe he and his work are great. This “cognitive dissonance” [14] is distressing enough have physiological correlates [15], and can lead us to alter the truth or our memories in order to allay it.

A system programmed to “feel proud” could face a similar situation. When a creation enters \mathcal{M}_i , the creator agrees with the critic’s evaluation, that is, $E_{ii}(c) \approx E_{jj}(c)$, which, if $\mathcal{E}_{ii} = \emptyset$, was arrived at via other critics’ preferences. When knowledge of these preferences or their weighting changes, some evaluations in \mathcal{M}_i could drop, as would M_i . By construction, the system cannot tolerate too large of a drop. Since it must also accurately represent others’ preferences, it cannot simply refuse to change that knowledge. To resolve this conflict, it could add information to \mathcal{E}_{ii} that keeps M_i from dropping too much.

False Inferences About Preferences Criticism includes the reasons behind an overall evaluation. However, the reasons we offer do not always reflect how we make our decisions. For instance, people will say why they preferred one of many products, all of which are actually identical [16]. Similarly, we invent reasons that sound good if our real reasons aren’t socially acceptable. For instance, though our evaluations of one aspect of a person pollute our evaluations of other aspects (the “halo effect”) [17], we often don’t know or admit this. Instead, we offer reasons that are demonstrably unrelated to our actual decision process [16].

A creator whose standards are solely based on other people’s standards is unlikely to say that she likes something “because he likes it too”. Instead, she will search for distinguishing features of the item to form a plausible-sounding explanation. Even if incomplete or incorrect, this utterance becomes part of how she understands her preferences, and might even impact future evaluations.

Suppose that a creative AI had a language for communicating criticism. Given that \mathcal{E}_i can consist of exemplars, neural networks, and other irregular representations, there is a large chance that the language could not express complete and correct information. If the rules it extrapolates are put into \mathcal{E}_{ii} , two things happen. First, the inaccuracy of the rules will lead to evaluations that no longer directly follow $\mathcal{E}_i \setminus \mathcal{E}_{ii}$. Second, the creator’s standards will lag behind changes in $\mathcal{E}_i \setminus \mathcal{E}_{ii}$, since those will not be reflected in \mathcal{E}_{ii} . Thus, the system will begin to develop divergent standards, albeit clumsily.

Selective Acceptance Seeking Even someone with a completely independent sense of what he likes might want a style somewhat similar to people he admires. If one such peer’s preferences shifted, he might adjust his own preferences in that direction. However, there would likely be several others peers he wishes to be somewhat near to, leading to experimentation until an equilibrium is reached.

Once a creative AI relies substantially on \mathcal{E}_{ii} , changes in other critics’ preferences will have a smaller impact on E_{ii} . However, it might try to keep an acceptably low discrepancy between $E_{ii}(c)$ and $E_{ij}(c)$, where j is a critic who i has a high affinity for. Indeed, this might be what enables the system to deviate from others’ standards but stay recognizably within the same domain or genre.

3.3 Autonomy Revisited

Creative autonomy requires autonomous evaluation, autonomous change, and non-randomness. The system could certainly be capable of autonomous and non-random evaluation. Furthermore, none of the schemes described above makes random changes at random times. Therefore, it just remains to be considered whether the system's changes would be autonomous.

In all three schemes, change happens in response to external events. For cognitive dissonance and acceptance seeking, this would be changes in others' standards ($\mathcal{E}_i \setminus \mathcal{E}_{ii}$), or in affinities (\mathcal{A}_i), possibly only after the effect of several separate changes accumulated. For false inferences, this would be the request for a critique. Unless another critic could manipulate when the creator starts and stops changing its standards, these change processes could be autonomous.

With only a single critic, such manipulation is possible. Acceptance seeking would simply entail following the lone critic's changes in standards within a margin of error. The critic could take advantage of false inferences by requesting criticisms as a way to perturb the creator's standards, only stopping when an acceptable result was reached. The critic could also give extreme and inconsistent ratings to trigger changes via cognitive dissonance.

The situation changes with multiple critics. The complex web of relations between \mathcal{A}_i , \mathcal{E}_i , and \mathcal{M}_i/M_i would make it hard to predict whether an external change would trigger adjustments to \mathcal{E}_{ii} . Multiple simultaneous changes might cancel out, or several small changes across time could accumulate until one change unleashes a string of compensatory adjustments. This would make the system less clearly responsive to any single critic, and in particular much more difficult for any single critic to manipulate.

Autonomy also precludes making fixed transformations to others' standards. When knowledge is first put into \mathcal{E}_{ii} , it is derived with error from others' standards, such as the stereotyped rules delivering criticism would produce, or the extrema that cognitive dissonance would enshrine. One could argue that this would only result in time-lagged caricatures of other critics' preferences. The counter-argument is that in a society where every creator pays attention to several other creators, such distortions would serve as attractors, leading to unpredictable clusters of similar styles, and unpredictable shifts in those styles. These emergent dynamics would make it impossible to say which creator was leading the change, meaning at least that no one creator was more autonomous than another.

This is obviously a strong claim that requires much empirical work. However, in a single-critic system, it is a fair guess that \mathcal{E}_{ii} would be more of a fun-house mirror reflection of the critic than anything that could be considered autonomously arrived at. Given that it would also be impossible to rule out that the critic had manipulated the system's dynamics until such time as it produced standards it liked, it seems fair to say that single-critic systems, at least as sketched here, would not achieve creative autonomy. The answer for multiple-critic systems will have to wait.

4 Conclusions

This paper introduced the concept of creative autonomy, which requires that a system be able to evaluate its creations without consulting others, that it be able to adjust how it makes these evaluations without being explicitly told when or how to do so, and that these processes not be purely random. A notation was developed to denote evaluations drawn from the integration of knowledge about several different critics' standards. Importantly, the system has different affinities for each critic, which impact how it integrates their opinions to form its own opinion. Initially it has no independently-held preferences, but this can change when it attempts to justify its evaluations, or if it must maintain high evaluations for some of its past work in the face of other critics' changing standards.

Such a system was argued to be capable of autonomous, non-random evaluation, and the change processes sketched are non-random. In a single critic society, it is unlikely that the system's standards would be more than distorted agglomerations of the critic's standards. What's more, the ease with which the critic could manipulate the system would make it hard to argue that the creator was changing autonomously. In a multiple-critic society, complex interactions might make any one creator impervious to manipulation, and emergent dynamics of the system could lead to clusters of creative styles. This awaits empirical demonstration, which would still leave open the philosophical question of whether changing absent direct manipulation is the same as autonomy.

The description of creative autonomy offered here captures only a small part of why humans can be considered creative. A system whose creations had a style that was not easily traced to a few influences, yet was still recognizably in same domain, would be a major accomplishment. However, as just mentioned, freedom from manipulation is not the same as acting purposefully. The system described here only makes changes in (possibly indirect) reaction to others' changes. Human creators, in contrast, proactively change their standards. It is conceivable that this system could make proactive changes by looking for patterns in how others' standards change with time and in relation to each other, which would be proactive, though perhaps not purposeful enough to be considered creative.

Though this work does not directly address the distinction between exploratory and transformative creativity, it could lead to interesting insights in that area. In particular, transforming a search space can be seen as searching over search spaces. In addition to making it indistinguishable from exploratory creativity, this begs the question of what objective function is used to select among spaces. Repeating this question over recursive searches of transform spaces of transform spaces of transform spaces, etc., one must ask what the base case is, i.e., what is the ultimate objective function? The perspective suggested here (and doubtless elsewhere, too) is that the ultimate objective function emerges out of the interactions between creators. It thus becomes essential for any system to be able to interact fully with its creative milieu if it is to be truly creative.

Creative artificial intelligence must always fight the impression that it is simply a fancy tool for expressing the programmer's creativity. This can lead to a desire to isolate the system from outside influences as much as possible. However,

as argued here, autonomy seems to require more, not less, interaction, though it must reach beyond the programmer. Though the hypotheses presented here are sketched with a broad brush and await verification, this work does suggest that creative AI must be viewed in a broader context than it traditionally has. Developing creative AI might still amount to solving an information processing problem, but a good part of this information comes from the social world. Of course, this is true of own creative processes, as well.

References

1. Boden, M.A.: *The Creative Mind: Myths and Mechanisms*. BasicBooks (1990)
2. Ritchie, G.: Some empirical criteria for attributing creativity to a computer program. *Minds and Machines* **17**(1) (2007) 67–99
3. Colton, S.: Creativity versus the perception of creativity in computational systems. In: *Creative Intelligent Systems: Papers from the AAAI Spring Symposium*, Stanford, CA, AAAI Press (Spring 2008) 14–20
4. Getzels, J.W., Csikszentmihalyi, M.: *The Creative Vision: A Longitudinal Study of Problem Finding in Art*. John Wiley & Sons, New York (1976)
5. Kasof, J.: Explaining creativity: The attributional perspective. *Creativity Research Journal* **8**(4) (1995) 311–366
6. Wiggins, G.A.: A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* **19** (2006) 449–458
7. Bown, O., Wiggins, G.A.: Modelling musical behaviour in a cultural-evolutionary system. In Gervás, P., Veale, T., Pease, A., eds.: *Proceedings of the IJCAI'05 Workshop on Computational Creativity*. (2005)
8. Jennings, K.E.: Adjusting the novelty thermostat: Courting creative success through judicious randomness. In: *Proceedings of the AAAI Spring Symposium on Creative Intelligent Systems*, Stanford, CA, AAAI Press (March 2008)
9. Nahemow, L., Lawton, M.: Similarity and propinquity in friendship formation. *Journal of Personality and Social Psychology* **32** (1975) 205–213
10. Kraut, R., Egido, C., Galegher, J.: Patterns of contact and communication in scientific research collaboration. In: *CSCW '88: Proceedings of the 1988 ACM conference on Computer-supported cooperative work*, New York, ACM (1988) 1–12
11. Cialdini, R.B.: *Influence: The Psychology of Persuasion*. Collins, New York (2007)
12. Zajonc, R.B.: Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology* **9**(2) (1968) 1–27
13. Mettee, D.R., Aronson, E.: Affective reactions to appraisal from others. In Huston, T.L., ed.: *Foundations of Interpersonal Attraction*. Academic Press, New York (1974) 235–83
14. Festinger, L.: *A theory of cognitive dissonance*. Stanford University Press, Stanford, CA (1957)
15. Croyle, R.T., Cooper, J.: Dissonance arousal: Physiological evidence. *Journal of Personality and Social Psychology* **45**(4) (1983) 782–791
16. Nisbett, R.E., Wilson, T.D.: Telling more than we can know: Verbal reports on mental processes. *Psychological Review* **84**(3) (1977) 231–259
17. Thorndike, E.L.: A constant error in psychological ratings. *Journal of Applied Psychology* **4**(1) (1920) 25–29