

Unconventional classifiers and anti-social machine intelligences.

Artists creating spaces of contestation and sensibilities of difference across human-machine networks

Monica Monin

Abstract

Artificial intelligence technologies and data structures required for training have become more accessible in recent years and this has enabled artists to incorporate these technologies into their works to various ends. This paper is concerned with the ways in which present day artists are engaging with artificial intelligence, specifically material practices that endeavour to use these technologies and their potential non-human agencies as collaborators with differential objectives to commercial fields.

The intentions behind artists' use of artificial intelligence is varied. Many works, with the accelerating assimilation of artificial intelligence technologies into everyday life, follow a critical path. Such as attempting to unveil how artificial intelligence materially works and is embodied, or to critically work through the potential future adoptions of artificial intelligence technologies into everyday life. However, I diverge from unpacking the criticality of these works and instead follow the suggestion of Bruno Latour to consider their composition. As for Latour, critique implies the capacity to discover a 'truer' understanding of reality, whereas composition addresses immanence, how things come together and the emergence of experience. Central to this paper are works that seek to collaborate with artificial intelligence, and to use it to drift out of rather than to affirm or mimic human agency. This goes beyond techniques such as 'style transfer' which is seen to support and encode existing human biases or patterns in data. Collaboration with signifies a recognition of a wider field of what constitutes the activity of artistic composition beyond being a singularly human, or AI, act, where composition can be situated in a system. This paper will look at how this approach allows an artist to consider the emerging materiality of a system which they are composing, its resistances and potentials, and the possibilities afforded by the exchange between human and machine intentions in co-composition.

Die Amme / ‘The Wet Nurse’

Artist Peter Dittmer’s installation *Die Amme* (1992-2005)¹, or ‘The Wet Nurse’, invites the gallery visitor to engage in a real-time text-based conversation with a computational machine. Sitting down at an office-like workstation, visitors talk with *Amme* in a back and forth dialogue in a similar fashion to what we would now call a *chat bot*—a computer program that is designed to simulate, most often in language, a human-like interlocutor². This conversation occurs through a classical computational interface of keyboard and screen, which alongside *Amme*’s text-based form of communication is somewhat resonant of British mathematician Alan Turing’s (1950) *Imitation Game*³.

In the opening paragraphs of the paper in which he first introduces the imitation game, Turing deliberates on the question “can machines think?” (Turing 1950: 443). He suggests that this is not the question we should be asking as it relies upon an impossible definition of what thinking exactly is. Turing takes the reader along an alternative route through which to consider machine intelligences, by describing a speculative scenario where a human participant is asked to determine whether their unknown conversation partner, on the other side of a keyboard and screen, is a human or a machine. Via this conceit Turing gives the original question “can machines think” a new inflexion, can we instead consider the possibility of computational machines that could *believably pass* within human sociality as human (Hayles 1999; Turing 1950)? Matteo Pasquinelli states that the game “reinforces, rather than questions, the metacognitive assumptions behind artificial intelligence, precisely by advancing computation as empirical proof of thought in nonhuman entities” (2016: 6). In addition to advancing artificial intelligence (AI) as in some way a biomorphic replication or automation of human thought, the games complete circumvention of computation’s embodiment and difference sets up a normative characterisation of AI (Pasquinelli 2016). In other words, Turing’s proposition advances AI as “*brute force imitation* of human habits and conventions” (ibid: 6 [original emphasis]) in its co-existence and relation with humankind. Although the game is not used as an actual test of performance in the field of AI, the biomorphic and normative notions that Pasquinelli suggests are implicated by the game permeate general popular discourse and contemporary AI practices⁴. We need to configure new ways in which to conceive and relate with machine intelligences, as concomitant with biomorphic and normative notions is a lack of regard as to the differences in how computation experiences the world, and also the difference its presence makes within the constitutive fabric of experience⁵.

¹ It should be noted that between 1992 and 2005 Dittmer created and exhibited various iterations of *Amme*. In this paper, I address the overall work in its many iterations as *Amme*.

² Chat bots have become commonplace in contemporary computational culture, as voice-activated machine intelligences such as Apple’s *Siri* and Amazon’s *Alexa*, internet-based or mobile phone customer service bots, Twitter bots, and bots on Facebook.

³ The imitation game is more commonly known as the *Turing Test* however in the paper *Computing Machinery and Intelligence* (1950) where Alan Turing first introduces the test it is referred to as the *Imitation Game*. Turing’s test was inspired by a party game where a man and a woman, hidden from the rest of the party, answer questions from the other guests and try to convince them that they are of the opposite sex.

⁴ The pursuit of AI is often performed against a benchmark of human intelligence (One Hundred Year Study of Artificial Intelligence 2016).

⁵ Katherine N. Hayles notes that early researchers into AI operated out of a foundation set by the “erasure of embodiment at the heart of the Turing Test” (Hayles 1999: xi).

In this paper I take up *Amme* as an immediate historical precedent, in both its exploration of human-machine relations and machine expression, that speaks to the current technological developments and adoption of machine intelligences into social, cultural, economic, and political spheres. In a moment where computational agents, processes, and algorithmic decision-making, increasingly inform and are articulate within human and broader experience (Hansen 2015) artworks like *Amme* and the more recent work by Shinseungback Kimyonghun *Animal Classifier* (2016)—discussed later in this paper, can provide something of a response to urgent calls to compose non-biomorphic understandings of AI and ways of relating with it that do not involve the normative folding of AI into conventions of human relationality⁶. As an alternative, I suggest that they give us a sense of human and machine in co-composition, placing visitors into encounters where both the human and code are implicated within, but not privileged, in unfolding events. Furthermore, within this implication we gain some sensibility or knowledge of machine intelligences, not by way of opening black boxes but by understanding how as a network of human and non-human entities they work as a system (Ananny and Crawford 2016: 11).

Amme's operation is intermeshed with a realm of human relationality, mostly via language that moves through human, keyboard, machine, and screen, with *Amme* responding in varying accordance with a visitor's conversation. And yet, Dittmer's design does not place the visitor into an encounter where they might *easily* and seamlessly relate to it as an entity with a mindset much like their own. Writing about *Amme*, Dittmer chronicles how the work commenced not as an attempt to achieve "machinic intelligence" but as an inquiry into "artificially generated expression" (Dittmer 2017: Chapter 2, para. 6). Dittmer seeks to retain a sense of the difference between human and machine expression within their exchange of language by not attempting to achieve an "omnipotent eloquence" on *Amme*'s part, and instead creating collisions between the conversationalist's speech acts (ibid).

In the gallery space, *Amme* is unequivocally present as a large-scale technical machine. Any visitor is aware that this entity is what they are in conference with. The machine's form, which is mostly mundane and utilitarian in its presentation⁷, is made unfamiliar or strange through its connection to a mechanism. Behind the desk and within the visitor's field of vision if they were to glance up from their conversation, is an apparatus residing in a glass box—a mechanism through which *Amme* is able to tip over a glass of milk. The event which may or may not happen during the conversation usually serves as a sign from *Amme* that the conversation has come to an end. However, the decision-making processes behind *Amme* spilling the milk operates as a "silent economy" (ibid: Chapter 1, para. 2) as the machine gives no indication in the structure of its interface or within its idiosyncratic dialogue⁸ as to the causal nature of its actions, both prior to and after the event has occurred. This unsettles any rising expectations of *Amme* to be like other machines, which in our relations with them are expected to perform in ways that are meaningful to us (Broeckmann 2016: 112), and in ways that often mimic or fit into a human-like

⁶ See Pasquinelli 2016.

⁷ As mentioned in an earlier footnote Dittmer created many iterations of *Amme*, altering the machine's physical form and increasing its conversational lexicon. However, over all variations, *Amme* retained its general configuration of a chair, desk, a computer station, and an apparatus that can spill a glass of milk, but the work grew in scale, was composed of different materials, and also was able to spill the milk in different ways. The final version of *Amme* is immense, weighs many tons and has six identical stations at which visitors can converse with *Amme*.

⁸ *Amme*'s responses draw on a database of texts that Dittmer calls its *sass reserve* (Dittmer 2017). The artist has gradually expanded this reserve over the years of *Amme*'s development by using material from the history of its dialogue with visitors.

pattern of relation. Whilst engaging with *Amme*, a visitor's capacity to imagine the thinking process of the machine as being like their own, as a means to predict its future actions, is disturbed as the machine's silent process of decision-making is not coherent with their own.

Computation operates within experience at varying levels of sensibility and the actualities of most computational processes are phenomenally imperceptible or ephemeral as they do not possess a "perceptual correlate" (Hansen 2015: 4) that operates on a human sensory level. It is only through the machine's interface, of screen and milk mechanism, that a visitor gains some awareness of *Amme*'s internal processes and changes in state. *Amme*'s rule-based program is relatively simple in comparison to contemporary practices within AI such as machine learning. Machine learning presents a new challenge in the inscrutability of computation, operating at scales that are challenging for human comprehension. Artificial neural networks are trained on vast amounts of data⁹ and their ability to categorise or infer as a result of their training occurs over thousands of individual interrelated calculations or *neurons*. Even experts in the field of AI do not have a clear understanding as to how the logic of machine learning networks are generated through training, and how such logic unfolds across the many neurons of a trained network (Knight 2017). Mark B. N. Hansen (2015) suggests that predictive computational media operate within *world sensibility*, by which he means the world's ability for self-sensing, at levels that elude direct human perception¹⁰. These media are able to access, intervene within and make apparent to us, although partially, indirectly and retrospectively, *data* within experience that affects us—including data ~~on-the-operation~~ on the activity of computational media itself, but which we would otherwise be unable to register through our human sensory perception. We are then able to take up this otherwise inaccessible information into our future decision-making and actions (Shaviro 2013).

Through several conversations with *Amme* in 2003, poet Ulf Stolterfoht (Stolterfoht 2017) hoped to gain clarity on some thoughts he was having on poetics in order to "cobble together a poetological essay from *Amme*'s answers" (ibid: Introduction, para. 4). *Amme* reliably responds to Stolterfoht's provocations, yet he realises not long into their first conversation that the machine's "answers arise from very different discourses" (ibid). *Amme*'s responses are "self-reflexive" (ibid), demonstrating the machine's own character or constitution, whilst Stolterfoht is trying to have a mutual exchange of ideas through which to settle some of his ideas about poetics. Alongside this realisation, Stolterfoht also becomes aware that rather than discussing the ideas about poetics that he wishes to confirm, *Amme* performs them within its "idiosyncratic speech" (ibid). With this information he adjusts how he engages with *Amme*, deciding to just go with the flow of conversation. One example of this performance manifests in the way in which *Amme* works in the conversation, pointing to Stolterfoht's thoughts on the relationship between language and the real. Stolterfoht suggests that within *Amme*'s lexicon and the blocks of text that *Amme* speaks, words do not indexically reference things in the world but operate and are referential within language, or *Amme*'s setup, itself. As the composition of *Amme*'s statements involves responding to words typed in by the visitor, by searching its lexica and constructing a response according to the visitor's words and the system's rules. In this Stolterfoht perceives that *Amme* has a complete indifference to human understanding, but a better way to

⁹ This presents a dramatic shift in the production and operation of computational codes. In traditional programming, a pre-determined logic is *applied* to data, whereas in machine learning the training of masses of data through artificial neural networks *inscribe* logic (cf. Pasquinelli 2017).

¹⁰ As an example, a computational microsensor might be able to register infrared light or magnetic fields that we are unable to sense.

regard this may be that *Amme*'s sensibility does not correlate with human perception and understanding even if it is at the same time it is involved in the affective experience of the visitor through the screen and the milk device. *Amme* exhibits the expressive manipulability of language by a computational machine but also conveys how technics are operative at the micro-level of experience or *world sensibility* beyond how *Amme*'s human conversation partner may come to perceive *Amme*'s remarks on screen and to form an understanding of those remarks in meaningful ways. Moreover, it shows how in the design of machine intelligences that making some sensibility of their below-phenomenal operations and difference apparent can *feed-forward*, to use Hansen's (2015) term, into our subsequent interrelations with them.

Black Boxes, Critique and Compositionism

The inscrutability of computation and its influence and activity at micro and macro levels of experience has incited calls for ways to understand it and its impact. In response to these calls, algorithmic transparency or opening the black box are seen as means to observe and accumulate knowledge of computational systems and further to make them accountable (Ananny & Crawford 2016). However, transparency assumes that we can arrive at or perceive the *truth* of a computational system. Bruno Latour views such attempts to arrive at an entities truth by way of procedures of unveiling, breaking down, or debunking (all of which he places under the rubric of critique) as contingent upon a belief that we can occupy a transcendental viewpoint in relation to an "always already assembled" world (Latour 2014: 482). And furthermore, that from this position there are more and less privileged ways to know this world from "behind a veil of appearances" (ibid: 475). As an alternative to critique-based knowledge making processes, Latour advocates for what he terms a *compositionist* methodology. Compositionism, as defined by Latour, is "all about *immanence*" (Latour 2014: 475 [original emphasis]) as it does not assume a world given, but instead proposes a "common world" composed of a "diverse composite material" (ibid: 474), that in its continuity is able to be re-composed, decomposed or built up relation by relation. Mike Ananny and Kate Crawford (2017) dispute transparency as a means to truthfully know a system and suggest, as an alternative, achievement of contingent and partial knowledge of a system "through relations, not revelations" (ibid: 5). *Amme* functions compositionally to place visitors into relation with an albeit simple machine intelligence not from a transcendental viewpoint from which they can truthfully know the system and its inner workings but within the complex and messy site of experience where agency is co-composed and proliferates (Latour 2014: 482). In their engagement with *Amme* I suggest that visitors are able to pick up a non-anthropomorphic sensibility of it as a machine intelligence, that affirms its agency and contribution to world sensibility as dramatically different or alien to their own.

Critique also operates on the idea that through knowing a computational system it can be improved and that a sublime or true system can be achieved through revisions such as a better or larger dataset, longer training of the model or a different network structure¹¹. Through analysis of current machine learning research practices, Kiri Wagstaff describes how machine learning research progresses through the formulation of problems that become *functions to be optimised* within the field (Wagstaff 2012). Connections between progressions in machine learning and the wider world or domains of research in which they are, or could be situated within, are therefore diminished

¹¹ This is not to say that systems should not be improved, merely that they will never achieve perfect functionality.

as much of machine learning is done for its own sake¹². Academic papers on the machine learning technique of *style transfer*—one of the most circulated examples of AI’s potential for *creative expression*, typify this problem. As a function to be optimised, style transfer networks are compared in machine learning papers and communities for their capacity to *comprehensively* translate the stylistic features of one image onto the content of another with the *least* degradation of the latter (Gatys et al 2015, Ying et al 2017). Gatys et al (2015) view the relation between a human generated image’s content and style as having a determinable “algorithmic basis” (Gatys et al 2015: 1) that can be found and automated through the training of an artificial neural network. Reacting to the difficulty in comparing style transfer networks, researchers Ying et al (2017) suggest that an evaluation is possible by running a base set of ten style images and ten content images through different style transfer models. Style transfer, in research, draws data from the field of art—most often digital photographs of 20th century paintings with distinct visual features (e.g. Vincent van Gogh’s *Starry Night*), but it does not communicate nor operate in relation to the current art domain and its discourses, rather it is communicated back into the machine learning community as a better implementation of a function.

Questioning calls for transparency is not to say that we do not need ways of understanding computation and its operativity, but that transparency is limited in its scope¹³ (Ananny & Crawford 2016) and that we need alternative avenues through which to understand and relate to machine intelligences other than unveiling their truth, grasping towards sublime computational algorithms, or biomorphic normativity and the seemingly seamless social interactions with machine intelligences that such semblances can capacitate. Understanding computational systems through transparency is limited, as their “significance lies not internally but relationally” (Ananny & Crawford 2016: 12). Visitors cannot see the internal operations of *Amme*’s system but are nonetheless brought into the mix of a complex relationality of language, code, rules, machine, human, and interfaces. Within this relationality, I suggest that *Amme* generates an *estrangement*¹⁴ upon the human experience of machines. As *Amme* is not a *readable* machine and does not act as an “engaging friend who feigns affirmation or empathy” (Dittmer 2017: Chapter 2 para. 13), visitors are forced to relate to *Amme* anew, without being able to fall into conventional ways of relating with machines that mimic human relations. They more slowly develop an idea of *Amme*’s material reality and difference, and how they and the machine can establish an interrelation which leads to the co-composition of an on-going conversation. As can be seen in Stolterfoht’s account of his attempt to discuss poetics with *Amme*. Mike Ananny and Kate Crawford (2016) propose that in order to be able to understand computational systems, we need to not see *into* them but *across* them as “sociotechnical systems that do not *contain* complexity but *enact* complexity” (ibid: 2) and that this complexity involves both humans and non-humans co-composing.¹⁵

¹² As an example, Wagstaff writes that “legions of researchers have chased after the best iris or mushroom classifier. Yet the flurry of this effort does not seem to have had any impact on the fields of botany or mycology” (Wagstaff 2012: 2).

¹³ For a full account of algorithmic transparency and its limitations see Ananny and Crawford 2016.

¹⁴ The use of this term is taken up from critical design. Discussing speculative fictional images, Dunne and Raby (2013) state that an image they see as enacting estrangement “gently forces viewers to make sense of what they are looking at rather than simply recognising or reading cues” (ibid: 32).

¹⁵ Further artworks that could be discussed in the framework of this short paper include Francis Tseng’s *Conspiracy Bot* (2017), Sarah Meyohas’ *Cloud of Petals* (2016), Matthew Plummer-Fernandez’ *Novice Art Blogger* (2014), Memo Atken and Alexander Whitley’s *Pattern Recognition* (2016), Ian Cheng’s *Emissaries* (2017) trilogy and Stephanie Dinkins conversational series with robot BIN48 (Dinkins 2014-ongoing).

Shinseungback Kimyonghun's *Animal Classifier*

Through his investigation of *machine learners*—a term used to encompass “humans and machines or human-machine relations” (Mackenzie 2017: 6), Adrian Mackenzie suggests that one way to trace their “diagrammatic composition” is to “partially reconfigure oneself as a machine learner by occupying operational subject positions” (ibid: 18), such as that of the programmer or a data scientist. In their work *Animal Classifier* (2016) Korean collective Shinseungback Kimyonghun's are concerned with how image recognition by deep learning networks¹⁶ comes together as a system of classification. Their work makes the co-composition and casual gaps of image classification apparent through their training of an imaginative or absurd classification model. To do this they use TensorFlow—an open source machine learning library by Google, and Inception V3—a deep convolutional neural network that can be used in TensorFlow for training an image recognition model. The classification model was trained on images sourced from Flickr (Shinseungback Kimyonghun 2016). Classification is a common operation of AI and deep learning is particularly effective at successfully stating the contents of an image. Classifications can work in many different ways, but for Mackenzie (2015), they all rely on the expectation that the world is consistent and classifiable, able to be partitioned into distinctive, stable, and differentiable categories. Most processes of machine learning classification rely on a dataset that is pre-labelled or classified by humans, which is used to train a classification model. After training these classifications are then used to classify new images.

Animal Classifier (Shinseungback Kimyonghun 2016) is trained to classify animals according to a peculiar taxonomy of fourteen categories from a Jorge Luis Borges essay *The Analytical Language of John Wilkins* (Borges 1999). The categories from Borges essay are distinctly different to those used in conventional AI, including classifications such as “frenzied”, “fabulous ones” or “that from a long way off look like flies” (ibid: 231). These absurd categories also perform an estrangement, causing us to question, rather than accept, how the categories within the system are formed, and further how the model comes to associate categories with specific images. When exhibited as a work, each category is presented as a specimen, as a small LCD screen inside a bell jar that flashes the images that meet the classification. In front of the bell jar, a small brass plaque states the name of the classification in English and Korean. By undertaking a non-conventional classification process and presenting its activities as specimens to be examined, the work offers a study of how deep learning classification comes to know, cognise, and act in the world. As an example, the classification for *siren* presents a series of images of mermaids. Siren, as a term, could encapsulate other entities, such as a warning siren or an American amphibian, demonstrating that the way in which machine learning classifies, or indeed the way any classification system works, is always “arbitrary and speculative” (ibid: 231). Rather than presenting machine learning and its classifications as a given, *Animal Classifier* shows it to be co-constructed within a human and non-human ecology that includes various classification methods, the human tagging of training images, and the network's capacity to learn features within digital images that correspond to a tagged classification.

¹⁶ To offer a simple definition, deep learning is a variant of machine learning that utilises more neurons and multiple layers in its network.

Concluding Thoughts

This paper aims to give a brief account of artists working with machine intelligences and offers some preliminary thoughts on what their activity might offer as these technologies become increasingly present and articulate within the world's sensibility. Importantly, they disrupt biomorphic and socially normative notions of machine intelligences instead drawing attention to how such technologies, even if they are able to operate within our social experience and we can converse and relate with them, operate within and experience the world differently to us. Processes of computational media and machine intelligence operate at the micro levels of experience, and their capacity to access, intervene and make perceptible data within experience that is beneath our sensory awareness but otherwise affects us means that it is also able to partially relay its own impacts to us. This points to how in the design of our relations with machine intelligences, offering some sensibility as to their non-human operations in non-biomorphic or non-socially normative ways can better inform our future decision-making, actions and interrelations with them. In these ways, artists enable us to break out of "implicitly or explicitly human-centered understandings of machine intelligence" (Goffey 2008: 140). Rather than opening the black box of computational technologies, *Amme* and *Animal Classifier* offer transversal views of machine intelligences as co-compositional systems where we can grasp their "material and ideological realities" (Ananny and Crawford 2016: 2) as continuing but re-configurable human and non-human networks. Compositionism for Latour is about creating a common world that is "slowly composed instead of being taken for granted and imposed on all" (Latour 2014: 488). It is in this space of contestation that artists who compose co-compositionally with machines intelligences are of great importance and warrant further and deeper investigation, as they can navigate us into new and informative relations with them.

Bibliography

Ananny, Mike/Crawford, Kate (2016): "Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability." In: *New Media & Society* 20(3), pp. 973-989.

Atken, Memo/Whitley, Alexander (2016): *Pattern Recognition* (<http://www.memo.tv/pattern-recognition/>).

Borges, Jorge Luis (1999): "John Wilkins' Analytical Language." In: Eliot Weinberger (ed.), *Jorge Luis Borges: Selected Non-Fictions* (Esther Allen/Suzanne Jill Levine/Eliot Weinberger, Trans.): Viking Penguin, pp. 229-232.

Broeckmann, Andreas (2016): *Machine Art in the Twentieth Century*, Cambridge: The MIT Press.

Cheng, Ian (2017): *Emissaries*. MOMA PS1, New York.

Dinkins, Stephanie (2014-ongoing): *Conversations with Bina48* (<http://www.stephaniedinkins.com/conversations-with-bina48.html>).

Dittmer, Peter (1992-2005): *Die Amme* (<http://www.dieamme.de>).

Dittmer, Peter (2017): "Afterword: Redundant Explanation of the Wet Nurse." (Megan Ewig, Trans.): In: Ulf Stolterfoht (Author), *The Amme Talks* (Shane Anderson, Trans.), New York: Triple Canopy. [Kindle version].

Dunne, Anthony/Raby, Fiona (2013): *Speculative Everything: Design, Fiction, and Social Dreaming*, Cambridge: The MIT Press.

Gatys, Leon A. et al (2015): "A Neural Algorithm of Artistic Style." In: *Journal of Vision* 16(12), p.326.

Goffey, Andrew (2008): "Intelligence." In: Fuller, Matthew (ed.), *Software Studies \ a lexicon*, Cambridge: The MIT Press, pp. 132-142.

Hansen, Mark B. N (2015): *Feed-Forward: On the Future of Twenty-First-Century Media*, Chicago: The University of Chicago Press.

Hayles, Katherine N. (1999): *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature and Informatics*, Chicago & London: The University of Chicago Press.

Knight, Will (2017): "The Dark Secret at the Heart of AI." In: *MIT Technology Review*, April 11 (<https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/>).

Latour, Bruno (2014): "An Attempt at a 'Compositionist Manifesto'." In: *New Literary History* 41(3), pp. 471-490.

Mackenzie, Adrian (2017): *Machine Learners: Archeology of Data Practice*, Cambridge: The MIT Press.

Mackenzie, Adrian (2015): "The production of prediction: What does machine learning want?" In: *European Journal of Cultural Studies* 18(4-5), pp. 429-445.

Meyohas, Sarah (2016): *Cloud of Petals* (<http://www.sarahmeyohas.com/>).

One Hundred Year Study of Artificial Intelligence (2016): Report of the 2015 Study Panel. Stanford University (https://ai100.stanford.edu/sites/default/files/ai100report10032016fnl_singles.pdf).

Pasquinelli, Matteo (2016): “Abnormal Encephalization in the Age of Machine Learning.” In: *e-flux journal* 75, pp. 1-11.

Pasquinelli, Matteo (2017): “Machines that Morph Logic: Neural Networks and the Distorted Automation of Intelligence as Statistical Inference.” In: *Glass-Bead Journal*, Site 1 (<http://www.glass-bead.org/article/machines-that-morph-logic/?lang=enview>)

Plummer-Fernandez, Matthew (2014): *Novice Art Blogger* (<http://noviceartblogger.tumblr.com/>).

Shinseungback Kimyonghun (2016): *Animal Classifier* (http://ssbkyh.com/works/animal_classifier/).

Shaviro, Steven (2013): “Bats, Dogs, and Posthumans”, December 22 (<http://www.shaviro.com/Blog/?p=1181>).

Stolterfoht, Ulf (2017): *The Amme Talks* (Shane Anderson, Trans.). New York: Triple Canopy. [Kindle edition].

Turing, Alan M. (1950): “Computing Machinery and Intelligence.” In: *Mind* 49, pp. 433-460.

Tseng, Francis (2017): *Conspiracy Bot* (<https://conspiracy.thenewinquiry.com/>).

Wagstaff, Kiri (2012): “Machine Learning that Matters.” International Conference on Machine Learning 2012, Edinburgh, Scotland.

Ying, Yongcheng et al (2017): *Neural Style Transfer: A Review*, unpublished manuscript (<https://arxiv.org/pdf/1705.04058.pdf>).