

Artificial Emotions in Virtual Storytelling

Interactive Storytelling is a notoriously challenging area, asking not only for a scientifically trans-disciplinary view, but also requiring to cut across and connect traditionally distinct communities of practice of researchers and engineers, established industries and professions, as well as enthusiastic new user groups. This dispersed and fragmented situation is mirrored in the range, growth and consolidation of international meetings and events of relevance over past two decades, with the International Conference on Interactive Digital Storytelling (ICIDS) traditionally representing a more encompassing, reflective view (including discussions on the very nature and purpose of digital storytelling), the title of the Interactive Narrative Technologies (INT) workshop series reflecting its technological origin and focus; and dedicated scientific, commercial, and user-group meetings (including the European Narratology Network (ENN), the AAAI's Artificial Intelligence and Interactive Digital Entertainment International Conference (AIIDE), and the Digital Games Research Association (DiGRA), covering assorted topical aspects. Against this background, the efforts of narratologists with affinity to technology (e.g., Janet H. Murray, Marie-Laure Ryan, and Nick Montfort); visionary technologists (e.g., Chris Crawford, Joseph Bates, Michael Mateas, and Noah Wardrip-Fruin); and topical projects (e.g., the seminar Oz project at CMU, and, more recently, the EU-funded Network of Excellence on Integrating Research on Interactive Storytelling: IRIS), as well as working meetings (including the Dagstuhl Seminar on Artificial and Computational Intelligence in Games, attended by the candidate in 2012) aim to develop coherent pictures and guidelines, identifying and reconciling key disciplinary contributions.

We found due consideration of this general setting and its (in particular: recent) historical development to be of importance for the appraisal of the submitted text in which (to freely paraphrase Valentino Braitenberg), the candidate can be construed to have followed a productive path integrating more prescriptive and engineering-oriented “downhill” invention with interpretative “uphill” analysis: Yet again, the availability and development of concrete, working tools and systems, and the pursuit of clearly grounded meaningful empirical goals in what has been termed the “Pasteur Quadrant” of research, have proven key enablers for tangible progress, in particular in such “wicked”—in the candidate's own and well-picked characterisation of the interactive storytelling domain—areas of research. *The thesis thus documents a productive tight integration of theoretical and empirical development.*

Overall, the thesis describes the candidate's scientific climb from a narrower notion of prescribed control and related “minimal engineering” towards support for broader behavioural emergence and resulting challenges of governance, and the development and evaluation of first related analytic tools, leading to considerations of increasingly “storytelling” and less “computer animation” nature: A particular noteworthy aspect we see with this well-framed and overall logically structured submission, lies in how it documents the development of the candidate's grounded *grasp* (here as favourable alternative to a mere intellectual logical “understanding”) of important challenges and issues in the field of interactive gameplay and storytelling, and of the *awareness* and *readiness* to appreciate the value of methodological and content-related contributions and critical discussions in the disciplinarily dispersed literature: These are *vital* requirements that *enable* and *motivate* active searching, a productive ability to recognise and appreciate the deeper implications of scientific findings, and productive engagement with the many different languages of the multifarious members of the interactive storytelling “motley crew” constituency. Clearly, over his extended effort **the candidate** has developed into the kind of “inter- and trans-disciplinarily ready” researcher called for in this field, who **meets the high demands** on capabilities of continuous self-orientation and **genuine ability for creative scientific work**.

It then is an almost inevitable concomitant characteristic of such *original creative* effort to leave room for improvement along many individual dimensions (e.g., in terms of (systematic) coverage of and grounding in relevant related work; scientific rigour of individual contributions; detail of documentation of the specific contributions¹; and depth of balanced scientific evaluation and discussion of the achievements), as also

¹ To some extent compensated by the submitted source code and running programs.

explicitly recognised by the author. I.e., in our evaluation we do subscribe to the value of a “scruffy” (cf. Robert P. Abelson, Roger Schank) *completing* of first-of-a-kind scientific journeys, of achieving a sufficiently complete picture of an undertaking, in order to enable an understanding of the interplay of the individual contributing parts and allow for the specification of requirements and performance profiles informed also by practical considerations regarding the task at hand (as opposed to purely abstract, theoretical considerations). And consequently we propose that what might also be viewed as a “collection of shortcomings” is rather to be understood as a valuable basis for motivated future scientific development: Next to the documentation of completed work, fostering such growth is a key value of PhD theses, particularly in cases such as the present one, where the text contributes more to opening up and furthering a path rather than providing a final word on some particular subject matter.

Next to these considerations, the thesis does provide a number of novel engineering and scientific results, in particular:

- A technological “minimal” (i.e., manageable, learnable, clean, efficient) toolset within the Pogamut framework, based on
- A related core set of theoretical conceptual abstractions, developed in particular against a broadly established operationalisation of practical reasoning (“BDI-model”);
- A furthering and grounding of the scientific discussion of how to scope and qualify emergent narratives, by demonstrating the application of a selection of theoretical measures to the concrete technological extended Pogamut framework that subserves the development and clarification of the *narrative* value of technical devices rooted in other domains (AI planning, philosophical practical reasoning).

Added to that, in order to demonstrate, fully integrate and materialise these contributions in an applicable way, and empirically evaluate them, the author further had to accomplish the non-trivial engineering tasks² of

- Connecting the Pogamut platform to the commercial Unreal Engine 2, in order to make a broader range of virtual settings readily accessible for scientific experimentation (as well as the evaluation of prototype systems and development of full-scale applications);
- Implementing two narrative games (SD One and SD Two).

These contributions are not only of particular importance for video game/storytelling systems engineering, but also for the membership of the interactive storytelling community at large, given that:

- The field is in dire need of *working* tools, for e.g. (and not limited to) the development of systems of non-trivial size (overcoming the frequent stranding of projects because of the multi-dimensional (theoretical, methodological, practical) complexities of developing a productive infrastructure); the enabling and support of scientific research in member disciplines (e.g., empirical grounding of literary studies or of psychological research); and for the inter-relating and comparing of results;
- They contribute to the conceptual furthering of the design of such systems, by providing a reference implementation of a specific set of core conceptual abstractions, and enabling a grounded evaluation and discussion, also against alternative proposals;
- They propose specific directions in which to further research in the field;
- The accessible recount of the author's scientific journey can be expected to serve as effective scientific scaffolding device not only for the formation of future students of similar background, but also, and importantly, for community members from other disciplines, enabling them to see some main opportunities and difficulties in the software engineering of interactive storytelling systems, opening up opportunities for fruitful collaborations.

Regarding formal aspects of the thesis, the text presents a clean overall structure: A general “Introduction” explicating the overall motivation and its refinement into a crisp set of four goals (p.7) is followed by chapters on “Theoretical Background”, covering related work; the “Emohawk Agent Architecture (EWA)”, into which the

² For these, the actual scope of the author's contribution is not clear, but while desirable, such information is not required for the present scientific evaluation task.

author's technological contributions have been integrated; the particular EWA-based “Applications” developed to empirically demonstrate the viability and value of EWA; and “Drama Analysis”, another significant theoretical and empirical challenge addressed in the thesis; and the “Conclusions” wrap up the results: i.e., how the four goals set in the introduction were all achieved, and which scientific and engineering contributions were developed by the author³. Complementing the “Bibliography” and six “Appendices”, the author also provided the content of a DVD providing the source code and installers for running versions of all developed applications, along with a collection of screen shots and videos of application runs.

The textual submission could have done with some more editing work (as only partly explicitly acknowledged by the author): At the macro/structural level, there are some singleton subsections; titles do not always present the best fit to the actual content—e.g., the chapter “1 Theoretical Background” in fact, and quite appropriately, also covers empirical work: it is (more) accurately characterised in the introduction to the second chapter as “introducing the Interactive Digital Storytelling field and the problems connected with developing a working IDS system” (p.35), this issue holds to some extent also for the very title of the thesis; the same related work is covered in different chapters, leading to some overlap, and to some local omissions; there are a number of *forward references*, where the reader has to know the content of later sections in order to understand the present text, and redundant content (e.g., Fig.3.7, p.78 vs. Fig.4.1, p.112). At the micro/local level, the particular selection of related work in the different chapters is not/ill motivated, creating the impression of a rather haphazard selection; the analysis of the different topical areas could be carried out in a more systematic fashion; the provided level of detail, in particular also of content covered in published work, but also e.g. in the legends of figures and tables, and the description of employed representations is less than what could be expected of a Ph.D. thesis: the reader should not have to resort to cited publications for basic information (no coherent overall description of the third-party MOSS system employed for evaluation purposes is provided), and the description of action representations is so scant as to not even touch upon issues such as canonical ordering of agent names in the representation of sub-scenes (section 4.4.2, p.108)—in contrast, a disproportionate amount of space is dedicated to describing Triangular Steering in Section 3.4.2 (pp.81-83), a third-party contribution integrated into the EWA architecture. Entries in the bibliography show assorted issues, e.g., of completeness, correctness, and coherent spelling/formatting; the text should have been spell-checked and proofed (in order to e.g. catch the (repeated) “casual” links on p.14) more carefully (and use of a grammatically competent correction tool would likely have contributed to raising the quality of language, which however poses no obstacle to acceptance in its present form).

Proposed aspects to be addressed and clarified in the Viva:

- *Believability*: The introduction establishes the development of “believable IVAs [Interactive Virtual Actors]” (p.7) as a core desideratum, corroborated by the narrow content of Section 1.2 (pp.16-18). This appears to be in contrast to the notion of “suspension of disbelief” (Samuel Taylor Coleridge, Janet H. Murray, Joe Bates, etc.) widely adopted in the interactive storytelling community (cf. e.g., Reilly, W. S. N. (1996). *Believable Social and Emotional Agents*. Ph.D. dissertation, School of Computer Science, Computer Science Department, Carnegie Mellon University, Technical Report CMU-CS-96-138).
- *Selection principles employed in surveys*: Please explain the principles underlying the selection of examples covered and evaluated in the sections, “Agent Architectures” (section 1.4, pp.21ff), “Overview of Current IDS Systems and Games” (section 1.5, pp.25ff), “Drama Analysis” (aka: story evaluation, section 4.1, pp.93ff), and “Action representation” (section 4.4.1, p.107f); and report on evaluation-relevant results and outcomes of the individual surveyed projects: These are covered only very selectively in these sections. (Note: This should also cover the unsubstantiated claim made in footnote 17 on p.33).
- *“Feeling”*: Please provide the motivation/discussion of how eqs. 2.1 (p.58) and 4.1 (p.103) for the calculation of “feeling” values match (which) intuitions and/or scientific findings (that the texts of Sections 2.7.3 (pp.56ff) and 4.3.1 (pp.102f) are lacking). As a corollary, explain also the relation of the notion of

³ Note that along with this summary clarification (p.129f), the clear documentation of the nature and the scope of the author's efforts throughout the text are essential information, given the integration of these contributions into the overall Pogamut platform and the resulting co-authorship of the cited peer-reviewed publications.

“feeling” introduced by you to the notion of “liking — how much is [sic!] the agent is attracted to another agent” (p.56) included in ALMA.⁴

As a corollary to this question, also relating to the use of conceptual terminology, clarify (justify) your idiosyncratic use of “*intention*” with respect to the established definition in the formal BDI model of Michael Bratman's theory of practical reasoning employed by the agent-oriented systems community.

- *Modelling: The origin/nature of constants*: Parameterisation of behaviour models is a well-known challenge, with the choice of specific values e.g. based on empirical findings, available performance models, or educated guesses: but it is important to always explicate the underlying rationale. Provide (a representative selection of) the missing explanations, e.g., for the “Feeling Implementation” discussed on pages 72ff, including with an analysis of the identified impact of your choices (by extension, this also regards the constants appearing section 4.4, e.g., subsections 4.4.4 and 4.4.5 on p.109, and in the item list of section 4.4.6 “Clustering”, on p.111).
- *Evaluation*:
 - The employed specification of “medium-sized drama” as “a computer game producing a coherent story that lasts for ten minutes” (Section 2, p.35)—i.e., the use of game duration as *only* criterion—requires justification. In particular, it should be contrasted to known published alternatives that employ additional conditions, such as frequency of *beats* (Mateas) or encountered *conflicts* (Szilas).
 - The statistical significance of the values reported in section 4.5 should be explicated.
 - We would appreciate a quick clarification whether the apparent clustering of blue and red values in the tension value plots (in particular, Fig.4.3, p.117), could be due to drawing artefacts (i.e., lines of one colour having been drawn after and over lines of a different colour).
- *Critical reflection: Towards “more narrative” tools*: With critical reflection on the developed “minimalist” contributions in a narratological context, how could the list of “requirements on the capabilities of the architecture” appearing on p.38 of section 2.1 “Problem Analysis”, be reformulated, so as to move away from the current set of permissive “*allowing for*” (i.e.: not hindering/preventing the user from) towards e.g., a stricter and more committing “*providing support of*”, as currently only stated in “R7. The architecture should support debugging of the agent behaviors”—if not in such overly generic manner?

In summary, Michal Bída has demonstrated his capability to identify, define and pursue a coherent line of independent research work to completion, resulting in original results that furthermore motivate further research. The amount of effort reported on in the thesis is substantial, and the technical difficulties of the engineering contributions must likewise not be under-estimated.

Therefore, **I do recommend, subject to a satisfactory performance at the Viva, that the candidate be awarded the title of doctor of philosophy (Ph.D.).**



Univ.-Lektor Dipl.-Ing.Dr. Paolo Petta

Austrian Research Institute for Artificial Intelligence
of the Austrian Society for Cybernetic Studies, Vienna, Austria (EU)
and
Dept. of Philosophy,
University of Vienna, Vienna, Austria (EU)

Vienna, 2017-09-11

⁴ You may also want to look into Elliott, C. and Siegle, G. (1993). Variables Influencing the Intensity of Simulated Affective States. In: John Horty and Yoav Shoham (Eds.): Reasoning about Mental States—Formal Theories and Applications. Papers from the 1993 Spring Symposium. Technical Report SS-93-05. Menlo Park, CA, USA: The AAAI Press. (pp. 58-67).