Proceedings of the 2013 Winter Simulation Conference
R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, eds.

SIMULATED COMPETITIONS TO AID TACTICAL SKILL ACQUISITION

Alexandre Feitosa

Departamento de Informática
Universidade Tecnológica F. do Paraná
Cornélio Procópio, PR 86.300-000, BRAZIL

Alexandre Direne, Fabiano Silva, Luis Bona

C3SL–Departamento de Informática
Universidade Federal do Paraná
Curitiba, PR 81.531-980, BRAZIL

Wilson da Silva

Fundação Cultural de Curitiba
Prefeitura Municipal de Curitiba
Curitiba, PR 80.230-040, BRAZIL

ABSTRACT

The paper presents a framework to support human skill acquisition of game tactics (*e.g.*, chess playing). We argue that cooperative work for defining heuristic constituents carried out by a learner should be alternated with simulated competitions to provide a formal, rating based feedback during the training phase. Firstly, our general definition of tactic concepts is bound to heuristic knowledge formalisations of two-player board games, including the notions of temporal, positional and material advantages. Secondly, our tactics definition language, aimed at the trainee, is described to cover a wide range of semantic features that can be applied in artificial games through a minimax search-based engine. The definition of heuristic parameters is based on variations of quantitative and qualitative production rules. The framework is instantiated by implemented software tools for the domain of chess. Finally, we draw conclusions about the suitability of the claims based on an empirical study.

1 INTRODUCTION

The purpose of this paper is to present a new framework of concepts and Artificial Intelligence (AI) tools to support human skill acquisition of game tactics such as the ones of chess or checkers. We also hope to emphasize a few interdisciplinary research interests lying beneath both computing and cognitive sciences to demonstrate educational benefits of on-line computer-based learning and coaching of heuristic games. The results of an empirical study support the main claims in that the problems encountered when trainees play a strictly heuristic game can be better solved if instead they externalise solely the tactic knowledge in formal terms for a machine to play on their behalf.

Indeed, the teaching of heuristic games, like chess and checkers, is traditionally carried out by initially presenting the learner with the basic rules for moving each piece. Instructors then move on to the intermediate-level of explanations by handling case studies where real board configurations are broken down into tactical components - as opposed to strategic ones (see a precise definition of tactics in Section 2). As a result, trainees develop novel skills that are useful for judging material, positional and temporal advantages. Material advantage refers to the relative piece forces whereas the positional one carries the positional configuration of the pieces, both estimated for a given stage of the game. The third class of advantage captures the evolving aspects of the two others in phased scopes, for rating the player as well as the opponent under equivalent quantitative measures.

Heuristic game abstractions are important because they tend to support problem solving capabilities in professional education (Sharples et al. 2002) but, although the adoption of tactical skills frequently leads to victory in a human-to-human game, there are several shortcomings with this approach. Firstly, many concepts about advantage are treated subjectively and indistinctly across the three referred classes of advantage. For instance, a well known tactics in chess is called board centre control, which should be accomplished soon after the opening moves. However, no formal representation is usually provided in traditional teaching as to when a game is precisely in the opening stage, which squares constitute the board centre and in what form a region can be considered controlled.

Secondly, even if learners assess game records that portray successful configurations, they will still find it difficult to identify and formally describe the precise individual advantages that the player envisioned when deciding on the heuristic plausibility of each move. This happens because the perception of tactical dimensions develops slowly over the years (Gott and Lesgold 2000), requiring the careful analysis of numerous complete games, alternated with actual human-to-human competition. In the past, many AI researchers claimed that there was little one could do to cut short this learning phase (Wagner 2012) without the help of new formalisms and software tools (Direne et al. 2009). In fact, real game practice seems to bear out such a claim since intermediate-level learners often show signs of over-general skills (Sharples and du Boulay 1988) that could be automatically detected but not easily bypassed (Bull and Kay 2007).

Thirdly, the discussion of learning skills solely based on actual competition is an imprecise process since, at the end of a game, players only know they won, lost or if it was a draw. In order to go a bit further, human instructors try to discuss some of the tactics applied by the learner through revisiting the game moves. However, both learner and instructor tend to manipulate different representations over the same content of typically procedural knowledge (Mozina et al. 2012), especially intuitive ones (Ainsworth 2006), which poses barriers for them to exchange meaningful explanations.

In short, the traditional pedagogic approach for developing tactical skills does not guide learners to understand the explicit meaning of heuristic factors as key-concepts. Nor does it encourage them to formalise, in a stepwise fashion, any decisions adopted in board analysis during a game (real or artificial). Such limitations are largely due to the lack of formal visual languages and software tools that allow experts as well as novices to precisely externalise (Munro et al. 1997) their intuitive assumptions about tactical advantage and, later, compare their effectiveness against other learners' formal tactics (e.g., through simulated competitions). In fact, language and tools for such purposes are the focus of this paper and constitute the dynamics of front-line research on computer science and human cognition to offer new modalities of skill acquisition (Hollan, Hutchins, and Weitzman 1987) through empirical testing of a learner's heuristic hypotheses (Estep 2006).

2 BASIC DEFINITIONS

This section clarifies our very specific notion of tactics in board games. It also goes through some of the specifics on how to operationalise cyclic comparisons among different tactics to support human skill acquisition.

2.1 Human Perception on Game Tactics

In both technical and common sense terms, the concept of tactical skill is normally vague, even in the context of two-player, board games. Thus, one fundamental factor in the point we make here is that it refers to the *short-term* visibility of the worths and effects taken into account when planning the next plausible move, as opposed to a longer-term one (*e.g.*, of strategic planning – not considered here). Typically, short-term visibility comprises only three or four combinations of moves ahead, which leads to replanning every turn during a game.

Due to the apparent limitation in the planning depth, it is crucial that the trainee player also perceives as early as possible the non-combinatorial side of a tactic. More specifically, it is important for the trainee to envisage any personal advantage as well as counter-advantage just by summing up his or her board constituents as well as the opponent's to check if they cancel each other out or not. For example, in chess, a board constituent would be "*possessing a queen*," which means safe material advantage only if the opponent does not have the same thing or an equivalent. Likewise, a queen is worth far more than a *pawn*, at least in the initial phase of a game. Material valuations in the chess literature frequently attribute 1 point to a pawn and then fill in the rest of the constituents with relative values of the pieces – say, 9 points to a *queen*, 5 to a *rook* and 3 to a *bishop* or a *knight*.

Besides material advantage, an effective tactic often includes definitions of positional and temporal constituents. An example of positional advantage would be a pawn closer to the opposite line from where it started on the board since it is more likely to be promoted, except under attack. In this case, the relative value of that pawn would be incremented, say, to three times the value of an ordinary pawn. An example of temporal constituent would be some sort of phase-transition rule which flags how near the end-game is (*e.g.*, by counting pieces on the board) and increases the value of every pawn, regardless of their position, because pawn promotion is the only way to recover material loss.

2.2 Constituents of Tactics as Heuristic Functions

Assuming that AI search engines already perform quite well for the combinatorial part of short-term planning, it makes sense to direct the trainee's attention mainly to the constituents of a tactic. By doing so, one should be able to work out and revamp such constituents using a domain-specific formal language. In our framework, a created setting can then be automatically interpreted and mapped directly into the heuristic function of a conventional *minimax* adversarial search engine (Fogel et al. 2005), freeing the trainee from the burden of manually carrying out the combinatorial process in every decision step for determining the next plausible move.

Thus, in our approach, the constituents of a tactic are nothing more than the components of a classical heuristic function that gets evaluated and returns a single numeric value. Moreover, such evaluation is done by taking just a board configuration as argument. This is relevant since the role played by the heuristic function in planning plausible moves very much depends on the precision of its component valuations. This will allow minimax search to perform further orderings and cut-offs among board combinations based on reliable precedence data.

2.3 Simulated Competition of Tactics

We define a simulated two-player game between two tactics as the application of the respective heuristics settings to the same minimax adversarial search engine, also preserving the depth for both game sub-trees (a hierarchical *And-Or* graph representation). Each activation is done in an alternate manner, by computing one plausible move from one tactic following one move from the other, until the game ends. This is in fact, no more than a comparison of a heuristic setting against another under the same combinatorial platform.

A simulated competition among more than two tactics is defined here as the combination of simulated two-player games according to some method of tournament matching (e.g., swiss and round-robin). Additionally, despite the alleged symmetry of certain games, there have been cases of famous players who showed preference for certain conditions. For example, in chess, one might prefer to play with white pieces whereas others with blacks. But the deterministic nature of heuristic functions does not leave much room for judging influences of the like. Therefore, even considering that matching methods take this issue partially into account, for a thorough assessment of tactic force, it is important to guarantee in simulated competitions that each pair of heuristic settings plays two games in a round: one for each piece colour or board side. By now, it becomes clear that an *effective* tactic is one that holds a good mark of the top rank in a simulated competition.

2.4 Cooperative Definition of Effective Tactics

Previous subsections concentrate on what an effective tactic is, but they do not address how one goes about writing one. No matter how accurate the tactic might be, it is widely known that heuristic values for expressing advantage do not have to work well in a hundred percent of the board configurations. Thus, in order to build an effective tactic, the trainee is expected to start by using his or her own empirical knowledge in an abstract reasoning way where the chosen heuristic components are of the essence. In our framework, each component can be described incrementally, either as a quantitative dimension (typical of material and positional advantages) or as a qualitative rule (typical of temporal representations), which also reverts to a numeric value.

But concocting such a setting is no easy task since it usually relies on a mathematical description of at least one's own perception on what is important in a board configuration. As a result, the trainee player should be prepared to face limits in representation and get help by cooperating with other trainees and masters (Caine and Cohen 2007). This can be realised by inspecting other trainees' settings and annotations, discussing alternative valuations and arrange to periodically take part in simulated games and competitions against other evolving tactics (all learner-centered tasks).

2.5 Alternating cooperation and competition

After an arbitrary number of simulated games against artificial or human opponents, if the trainee doesn't manage to improve the heuristic setting, it will be curtains for the tactic as a whole. After that, a new setting will eventually replace the old one through some cooperative, educational cycle. This suggests more study and more cooperation which, in turn, eventually leads to a new need for competition, closing a cycle that alternates cooperation and competition, indefinitely.

Ultimately, the present work takes a step further and argues for the coordinated alternance of the cooperation and competition phases. According to recent findings in education, when conducted in a supervised manner, skill learning can benefit from metacognitive processes depending on whether feedback is realistic or not (Power and Scott 1998). The way our competition phase is wired to cooperation naturally engages the trainee in reflexive tasks. In doing that repeatedly and compassed by the coordinator, one is more likely to detect flaws in the valuation scheme of the heuristic setting. The natural next step then is for the trainee to react accordingly and revamp the setting by incrementing newly learned components from the cooperation phase and balancing them along with the existing ones.

3 LANGUAGE AND SOFTWARE TOOLS

Our framework consists of a language and software tools for managing various tasks of heuristic board games. It includes a domain-independent model of interpretation, integrated with the tools, for controlling on-line social interactions among users and carrying out both human-to-human and simulated competitions. Although a large proportion of the current research effort has been dedicated to conceive language design principles and to implement their related software tools, for brevity reasons this section only glimpses the theme among the other issues just before getting into practical application.

3.1 The DHJOG Language

The *DHJOG* language is aimed at allowing trainees and game practitioners to formally represent the non-combinatorial part of a tactic, namely the heuristic setting. On the one hand, DHJOG is constrained to the domains of two-player, heuristic board games. On the other hand, it is quite general in the sense that its visual and textual structures can be used do capture the estimates based on diverse board and piece configuration features. As long as basic arithmetical, geometrical and algebraic aspects are concerned, it can be viewed as an open-ended formalism for instructing artificial players as to what is advantageous in material, positional and temporal terms.

As part of the language structure, DHJOG primitives are grouped by what we call *levels of expressivity*. The higher the level, the more skilled the trainee has to be in order to make use of its primitives. A list of brief descriptions of the language primitive functionalities per level number follows:

1. as the most basic level, it allows only relative valuations of the pieces to express material advantage;
2. positional advantage is allowed through the visual definition of geometric board features;
3. global board valuations can increment previously defined heuristic components;
4. phase compartments as well as phase-transition rules can be added to express more specific valuations in different game stages (e.g., opening, mid or end-game);
5. Provides a wider range of text-based primitives to represent conditional expressions of phase-transition rules;
6. allows the extension of language primitives and interpretation procedures of the underlying engine (i.e., by directly programming Java code).

It is still worth mentioning that code in DHJOG tends to be compact because in the large majority of the situations, the language dispenses with the need to specify opponent equivalents, unless they do not exist or when a symmetric description has to be added (in either case, action must be taken by the trainee). An example of the latter case would be in related to *kingside* or *queenside* castling in chess. Whatever the option, the symmetric board region in front of the opponent's king final position has to be geometrically specified if any incremental piece valuation context refers to it (*i.e.*, the symmetric region cannot be determined automatically). Likewise, an example of the former case would be to explicitly prevent incremental piece valuation of one colour to happen within the boundaries of a region. Finally, it does not harm to insist that any heuristic setting defined in DHJOG will be interpreted by the related software tools as a function that returns a numeric value for ordering purposes of minimax search.

3.2 The HeuChess Software Tools

An instantiation of the framework, named *HeuChess*, was implemented as a cloud-based software tool for the domain of chess. It allows the user to interactively describe a heuristic setting in a dialect of DHJOG. The software interface and server tools were implemented in Java. They allow identified access to the system, although shared data is mandatory through the tournament supervisor instructions. This permits putting into practice the alternation of cooperative and competitive phases. The web address of the main page is http://www.alexandrefeitosa.com.br/heuchess.

On the interface side, trainees can easily create heuristic settings by directly manipulating graphical and textual objects. Inspection of other users' settings is also permitted during the cooperative period. Figure 1 shows a snapshot of the interface set under level-2 primitives for defining new piece valuations which will be used later when a mid-game phase is created under level four. Notice that a board region was added as part of a positional advantage scheme that attempts to formalise the well-known tactic concept of board centre control. All interfaces are in Brazilian Portuguese.

Due to the flexibility of HeuChess underlying code interpreter, even the very first components of a setting can have their tactic force quickly inspected by the user, which can be realised with the tools in three ways. One way is by functionally computing and showing the actual numeric value of a specific board configuration entered manually. Another is when the user plays a game directly against a minimax-based robot coupled with the heuristic setting as its board evaluation function. The third way is by running simulated competitions of the currently defined setting against other settings. For that, HeuChess allows certain competition features to be set, such as the matching type (*e.g.*, round-robin), game tree depth, scoring scheme and others. Figure 2 shows a snapshot of the tournament management tool showing competition progress data, game tree depth and matching features.

In face of bad performance under any of the three assessment methods, the tools allow the trainee to squander each component of the setting separately. Such a procedure starts with the learner selecting from

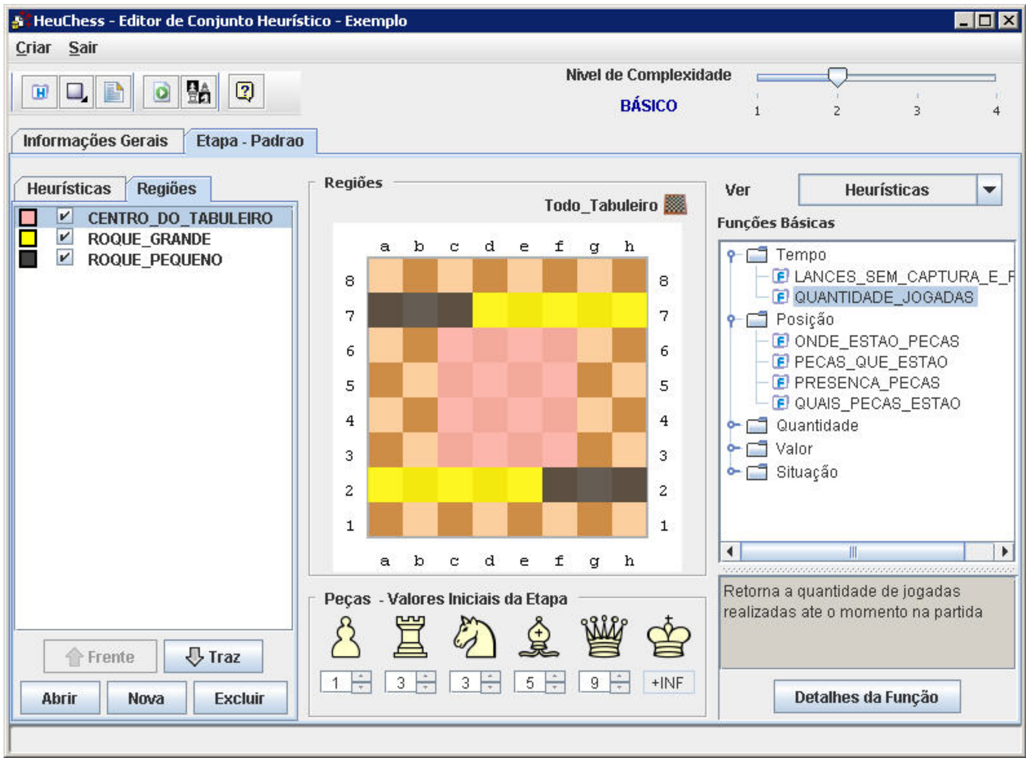

Figure 1: HeuChess interface set under level-2 primitives.

his or her own gallery (it is important to have one) a heuristic setting collected in the past. It must be good enough to play a challenging simulated game against the setting which is being tested. Under the squandering mode, HeuChess allows stepwise inspections of each component value of the tactic as well as graphical game subtrees, including their minimax values, from best and worst moves. Figure 3 shows a snapshot of the tool working under this mode. Like in debugging tools of programming languages, note the many interface items being inspected, including the current board configuration also in algebraic notations, such as the Forsyth-Edwards one.

After imprecisions and errors are found, incremental definitions can be entered and reassessed. This can be done in parallel with the inspection of other trainees tactic codes and by educational discussions. In the normal pace of events, the competition supervisor (*e.g.*, the course instructor of a class) can use privileged access to start a new tournament to which trainees can register their best heuristic settings. After the timetable is defined, public access to such settings is suspended until competition ends.

On the server side, HeuChess keeps detailed records of users' entries. Both tactic code and access data captured from the interactions during the cooperative lessons are kept in its database. Access permission data is also handled by privileged users. The current service is part of a wider framework from which other project tools, data repositories and open-source code are available. The main project site is accessible at http://xadrezlivre.c3sl.ufpr.br/projeto.

4 EMPIRICAL STUDY

In our investigation, we choose the domains of chess tactics, particularly the ones that focus on mathematical knowledge (*e.g.*, arithmetical, geometrical and algebraic). Tactical decision in chess is a hard task, especially when combined with the abstract abilities of formal language use required for the definition of detailed game heuristics. Despite this, expert chess players (unlike novice trainees) are known for accurately identifying game situations and for justifying their empirical assessment in the form of heuristic settings. Thus, the

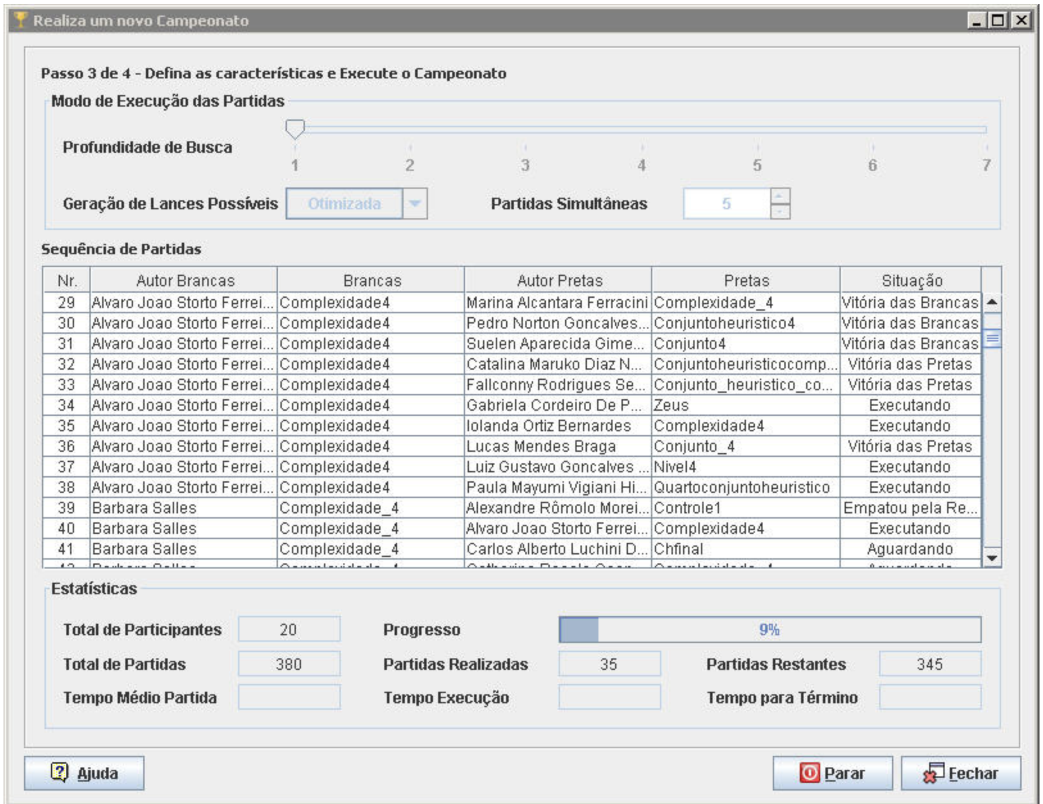

Figure 2: Screen showing competition progress data and matching features.

study with chess masters and their trainees was carried out in a college of technology where chess has long been part of the curriculum.

4.1 Aims

The aim of the study was threefold:

- to connect trainees via the Internet through our implemented software tools that allow for the definition and mutual exchange of chess tactics as well as the simulated competition among such tactics;
- to create a repository of elicited tactical definitions interspersed with social interaction data produced cooperatively and competitively;
- to determine any relationships that reveal the potential of human cooperation in developing tactical skills considered effective solely based on the capability to win simulated competitions.

Briefly, these were preceded by a series of formal training sessions conducted by human tutors. The design of the study was informed by extensive consultations with chess masters to determine the scope of tactical knowledge to train on, the most appropriate coaching approach to use and the subjects to be involved. Coaching was done according to a wide degree of freedom in terms of means and formats (e.g., use of real chess boards and text books (Salen and Zimmerman 2004)) that, in addition to the DHJOG language, were utilised during the sessions.

Figure 3: Tool working under squandering mode of component values.

4.2 Methodology

The study involved 88 college students, typically seventeen year olds, as subjects: 32 in their first-year and 56 second-years. All of them were in the novice to intermediate-level range as trainee chess players. Beyond the main claim of this research, the slight difference in age tried to capture the supporting idea that abstract mathematical skills were also important and thus should not vary too much across subjects when making extensive use of DHJOG's syntactic and semantic aspects.

Chess lessons were conducted as group discussions on a classroom-like scheme and as individual dialogues focussing on one trainee at a time. A bottom-up coaching approach was often adopted by allowing the trainees to begin with their own hypothesis based on a specific game feature. In chess textbooks, authors normally approach the reader in a top-down fashion (Guid et al. 2009). Despite this, we chose the former approach since our focus here was not on styles, and since a bottom-up approach is more likely to reveal the reasoning behind the trainees' judgements. As a supplementary coaching directive, the tutor was asked to interfere fairly often to make trainees externalise their reasoning chain. This was expected to give trainees a better chance to exhibit the mathematical aspects of their tactics and thus link them to the DHJOG language.

After a period of two months of formal lessons and cooperative use of the software tools, a final and more directed part of the study comprised the creation and competition of trainees' tactics. In order to highlight primarily the novice strengths, we applied problem statements that demanded a fair amount of tactical skills, while keeping empirical ideas to an intuitive context. Besides, case problems required the first four of DHJOG's levels of expressivity (presented in Section 3) to be understood and solved. A sample problem statement is as follows: "Write a chess tactic based on heuristic settings that cover levels from 1 to 4. The settings must include a specific valuation for each phase of a game, including a Four Knights opening. For the mid-game, the tactic must try to favor board centre control and, closer to the end-game, it must lead to pawn promotion."

4.3 Data Collection

The ultimate goal was for each trainee to write a tactic in DHJOG as a solution of each problem - a task that novices often failed to accomplish. Only 19 trainees managed to fulfill all the solutions using level-4 language primitives. They all managed to use language resources such as the phase menus, phase-transition rules, geometric board regions, incremental piece and board valuation. We collected the coded tactics created by these 19 trainees and carried out a comprehensive visual analysis over their details.

After that, simulated competitions, based on minimax search over game trees of depths 1, 2 and 3 were run to compare all the solutions of each problem. We recorded detailed logs of the round-robin matchings as well of as game moves and results where each matched round was rated by two games (different piece colours for each artificial player). These were then combined with the cooperation data logged during the cooperative training sessions in order to identify evidence of human learning about effective tactical skills. A full description of the collected data is presented in a technical report (Feitosa 2006). A summarized discussion on the data follows.

4.4 Discussion on Skill Acquisition

Because of the hermetic nature of most existing computer-based tutorials for games, they only permit trainees to learn heuristic relationships through competition, rather than explicitly discussing these relationships with others during the interactional periods. However, knowledge acquired by induction is fragile (Direne et al. 2008) and, as a result, learners could be expected to show misconceptions even after using such systems. One source of heuristic error we found in our study is based upon the possibility of learners assuming facts about board configuration which is in accordance with purely aesthetic mathematical concepts, but in practice do not occur. That is, as instruction proceeds, trainees' beliefs are expected to become over- and under-general, needing cooperative interventions to avoid the problem.

Table 1: Cooperation and tactic code data of the top twelve trainees.

C1	C2	C3	C4	C5	C6
1	0.909	9.6	10	26	1
2	0.8976	9.1	12	104	4
3	0.86	8.1	18	26	3
4	0.8783	9.8	13	117	15
5	0.8583	8.5	12	26	7
6	0.8552	8.4	20	52	10
7	0.8321	7.8	12	91	3
8	0.9476	9.6	15	0	6
9	0.74	7.0	15	91	0
10	0.8159	7.3	20	0	3
11	0.9438	10.0	21	0	0
12	0	0.0	0	0	0

The meaning of each column code in Table 1 is:

C1: final position in simulated tournament rank;
C2: average grade in college academic records (0.9 means 90%);
C3: maths grade in college;

C4: quality of tactic components defined in heuristic setting;
C5: number of cooperative accesses to other trainees tactic definitions;
C6: number of simulated games manually activated against other trainees' tactics comprising either level-3 or level-4 primitives.

To substantiate such claims, what the data in Table 1 seem to suggest are as follows:

- there is quite a direct relationship between maths skills and effectiveness of the tactic;
- exceptions are in ranks 8 and 11, where no cooperative access to other trainees tactics have been recorded by the HeuChess server;
- the number of tactic components of the heuristic settings is almost the same for all top 12 trainees;
- trainees holding the worst positions in rank are those who did not make cooperative access or did not run games manually to test their tactic definitions.

5 CONCLUSION AND FUTURE RESEARCH

The framework outlined provides a visual and textual language for the cooperative development of human skills on board game tactics integrated with a general model of heuristic search evaluation and simulated competition. Educational and computational aspects of the method and software tools have been discussed, including solutions for specific problems like dealing with the trainee's perception of short-term effects of heuristic valuations on the effectiveness of an automated tactic. The cooperation and competition prototype tools are implemented and allow instantiations of their shells to a range of domain-specific heuristic games.

To substantiate generality claims about the framework and the DHJOG language, we combined the evaluation of the software tools with their instantiation to the domain of chess. The empirical study concentrated on the potential of human-to-human cooperation as well as in simulated competitions to reveal the potential of combining both in alternating periods. The positive results obtained so far suggest the suitability of the framework for the development of novel training methods for tactical skills since heuristics is important for many areas of human creativity, ranging from arts to engineering.

Future research concentrates in adding new features to the framework in two ways. Firstly, we are working in a deeper integration of interface and server tools. On the interface side, more parameters will be available to improve cooperation tasks as well to set and monitor tournament conditions. On the server side, we plan to implement parallel processing of the competitions so that machine-demanding modules can return their results faster to trainees and coaches.

Secondly, the effort is in extending the DHJOG language with quantitative temporal concepts is underway. In this sense, game time would be a DHJOG primitive computed automatically as a strictly numeric function that returns some measure in machine clock units. This would then be applied by trainees in their tactic description to inform the minimax robot to halt search earlier than a given depth limit based on a certain amount of time dynamically calculated as a heuristic value. As a result, a brand new path of semantic relevance would open to trainees since actual game practice, on-line or otherwise, is all carried out under timed plays. Under such conditions, the robot will now have return the next plausible move based on result of the minimax search previously completed one depth shorter than that when halting occurred. It is important to say that our current general-purpose competition machinery is incapable of dealing with this new search mode as it requires an iterative-deepening algorithm (Sadikov and Bratko 2008) to be coupled to minimax. But that is not the main problem here since it is the educational context that matters most. Hopefully, continued AI work on computing and cognitive sciences will help in bringing about new languages and software tools for similar purposes.

ACKNOWLEDGMENTS

This research and development project was partly sponsored by grant number 20121015152542028–FNDE/MEC/Brazil, aimed at funding specific activities of the phase called UCA (Um Computador por

Aluno). We would also like to thank the two main Brazilian national research councils (CAPES and CNPq) for their support in different phases of the work.

REFERENCES

Ainsworth, S. 2006. "DeFT: A conceptual framework for considering learning with multiple representations". *Learning and Instruction* 16 (3): 183–198.

Bull, S., and J. Kay. 2007. "Student Models that Invite the Learner In: The SMILI Open Learner Modelling Framework". *International Journal of Artificial Intelligence in Education* 17 (2): 89–120.

Caine, A., and R. Cohen. 2007. "Tutoring an Entire Game with Dynamic Strategy Graphs: The Mixed-Initiative Sudoku Tutor". *Journal of Computers* 2 (1): 20–32.

Direne, A., D. Scott, F. Silva, L. Bona, M. Castilho, M. Sunye, and L. Garcia. 2008. "Acquiring expertise in medical radiology through long-term interactions". In *Proceedings of the 21th IEEE International Symposium on Computer-Based Medical Systems (CBMS2008)*, 403–408.

Direne, A., D. Scott, F. Silva, L. Bona, M. Castilho, M. Sunye, and L. Garcia. 2009. "Authoring adaptive tutoring systems for complex visual skills". In *Proceedings of the Nineth IEEE International Conference on Advanced Learning Technologies (ICALT-2009)*, 160–164.

Estep, M. (Ed.) 2006. *Self-Organizing Natural Intelligence: Issues of Knowing, Meaning, and Complexity*. New York: Springer.

Feitosa, A. R. M. 2006. "Definicao Formal de Taticas de Xadrez por Meio da Autoria Incremental de Conceitos Heuristicos". Technical Report PPGINF-191, dissertacao de mestrado – Departamento de Informatica, Universidade Federal do Parana (UFPR).

Fogel, D., T. Hays, S. Hahn, and J. Quon. 2005, April. "Further Evolution of a Self-Learning Chess Program". In *Proceedings of the IEEE CIG-2005 Symposium on Computational Intelligence and Games*, 73–76. Essex, UK: IEEE.

Gott, P., and A. Lesgold. 2000. "Competence in the workplace: How cognitive performance models and situated instruction can accelerate skill acquisition". In *Advances in instructional Psychology*, edited by R. Glaser. Lawrence Erlbaum.

Guid, M., M. Mozina, A. Sadikov, and I. Bratko. 2009. "Deriving Concepts and Strategies from Chess Tablebases". In *ACG-09: Proceedings of the 12th international conference on Advances in Computer Games*, 195–207.

Hollan, J. D., E. L. Hutchins, and L. Weitzman. 1987. "STEAMER: An Interactive, Inspectable, Simulation-Based Training System". In *Artificial Intelligence and Instruction*, edited by G. Kearsley. Addison-Wesley.

Mozina, M., M. Guid, A. Sadikov, V. Groznik, and I. Bratko. 2012. "Goal-Oriented conceptualization of procedural knowledge". In *ITS-12: Proceedings of the 11th international conference on Intelligent Tutoring Systems*, 286–291.

Munro, A., M. C. Johnson, Q. A. Pizzini, D. S. Surmon, D. M. Towne, and J. L. Wogulis. 1997. "Authoring Simulation-Centered Tutors with RIDES". *International Journal of Artificial Intelligence in Education* 8:284–316.

Power, R., and D. Scott. 1998. "Multilingual authoring using feedback texts". In *Proceedings of the 17th International Conference on Computational Linguistics (COLING/ACL-98)*, 1053–1059.

Sadikov, A., and I. Bratko. 2008. "LRTA* Works Much Better with Pessimistic Heuristics". In *ECAI-08: Proceedings of the 18th European Conference on Artificial Intelligence*, 897–898.

Salen, K., and E. Zimmerman. (Eds.) 2004. *Rules of play - Game design fundamentals*. Cambridge, Massachusetts: MIT Press.

Sharples, M., and B. du Boulay. 1988. "Knowledge Representation, Teaching Strategy and Simplifying Assumptions for a Concept Tutoring System". In *Proceedings of European Conference on Artificial Intelligence (ECAI-88)*, 268–270.

Sharples, M., N. Jeffery, B. du Boulay, B. Teather, D. Teather, and G. du Boulay. 2002. "Socio-cognitive engineering: a methodology for the design of human-centred technology". *European Journal of Operational Research* 136 (2): 310–323.

Wagner, G. 2012. "SIMURENA: A Web Portal for Open Educational Simulation". In *Proceedings of the Winter Simulation Conference*, edited by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. Uhrmacher, 1–12. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

ALEXANDRE DIRENE received his PhD in Computing in 1993 from Sussex University, UK. He is a Senior Researcher and Lecturer in Artificial Intelligence and a former Head of the Computing Science Department at the Federal University of Paraná (UFPR). His main research interest lies in the application of Artificial Intelligence techniques to educational software such as the design and implementation of tutoring systems for heuristic games. He has served as program committee member and reviewer of many national and international conferences and journals as well as Chair of several national conferences and workshops, including the *Special Interest Group on Computers in Education* of the Brazilian Computing Society. He is on the editorial board of the *Brazilian Journal of Computers in Education*. His email address is alexd@inf.ufpr.br and his web page is http://www.inf.ufpr.br/alexd.

ALEXANDRE FEITOSA received his MSc in Computing in 2005 and is currently enrolled in a PhD course in Computer Science at the Federal University of Paraná. He is a Lecturer in the Department of Informatics at Federal Technological University of Paraná (UTFPR). His areas of research interest include Interface Design, Game Development, application of Artificial Intelligence and the use of IT in Education. His email address is alexandrefeitosa@utfpr.edu.br and his web site is http://www.alexandrefeitosa.com.br.

FABIANO SILVA received his PhD from Federal University of Technology in Paraná (UTFPR), Brazil. He is a Researcher and Lecturer at the Federal University of Paraná and current deputy head of the Department of Informatics. He has been working on Planning and Optimization intelligent searching techniques for more than a decade now. His research interests include applying the techniques to simulated games. He is a member of the Brazilian Computer Society. His email address is fabiano@inf.ufpr.br and his web page is http://www.inf.ufpr.br/fabiano.

LUIS BONA received his PhD from Federal University of Technology of Paraná (UTFPR), Brazil. He is a Researcher and Lecturer at the Federal University of Paraná and a former head of the Department of Informatics from 2008 to 2012. He has been working on large research projects in Distributed Systems since year 2000. His research interests include Distributed Algorithms, Cloud Computing, Digital Preservation and Open Source Software. He is a member of the Brazilian Computer Society and the IEEE. His email address is bona@inf.ufpr.br and his web page is http://www.inf.ufpr.br/bona.

WILSON DA SILVA received his PhD in Education from Estate University of Campinas (UNICAMP), Brazil. He is a Researcher of Fundação Cultural de Curitiba and a Lecturer at SION. His publications approach human acquisition of concepts through interactive simulations by dealing with learning objects. He has been working with new techniques for chess coaching, including those based on features of grand-masters' expertise. His email address is wilsilva@onda.com.br and his web page is http://www.wilsondasilva.com.br.