In the Proceedings of the 2021 AAAI Fall Symposium "Cognitive Systems for Anticipatory Thinking - 3rd Wave Autonomy", Arlington, VA, November 4-6

A Framework for Deep Anticipatory Intelligence Analysis

Gheorghe Tecuci, Dorin Marcu

Learning Agents Center and Department of Computer Science, School of Computing, George Mason University, Fairfax, VA 22030, USA
tecuci@gmu.edu; dmarcu@gmu.edu; http://lac.gmu.edu

Introduction

Anticipatory intelligence is the complex task of identifying and assessing new, emerging trends, changing conditions, and underappreciated developments to challenge long-standing assumptions, encourage new perspectives, identify new opportunities, and provide warning of threats to national interests, based on current information of all kinds that come from a variety of different sources (ODNI, 2019). It is aimed at potential events, including low-probability high threat events and includes active attention management, focusing attention on likely sources of critical information (Klein et al., 2007).

The prevailing approach to anticipatory intelligence analysis, and intelligence analysis in general, is a *holistic* one where the analysts, after reviewing large amounts of information and mentally processing the data, reach a conclusion (Marrin, 2011). Consequently, there is very limited *transparency* on how exactly the conclusion has been reached from evidence, what assumptions have been made, and how exactly the probability of the conclusion and the confidence in this probability have been assessed.

Complementary to the holistic approach is the *structured analysis* approach, but the current practice relies on very simple analytic techniques, such as those described by Heuer and Pherson (2011). The highly acclaimed book *"Critical Thinking for Strategic Intelligence"* (Pherson and Pherson, 2021) only presents general guidelines for good analysis, but there is no clear process on how to assess the probability of a hypothesis and the confidence in that probability. Much more advanced methods use Bayesian probabilistic inference networks but, despite their implementation in advanced analytical tools such as Netica (https://www.norsys.com/), developing a Bayesian network is a very difficult task for an intelligence analyst.

Advances in sensor technology have made it very easy to automatically collect massive amount of data, leading to the so-called *"deluge of data."* This huge and growing gap between the ability to collect information and that of analyzing it, is making increasingly difficult to draw *timely* and *accurate* conclusions from these *massive* amounts of collected data, much of which is irrelevant to the analysis.

This paper presents a framework for anticipatory intelligence analysis to be implemented into an instructable agent that can be taught how to assist an analyst. It is based on a series of previously developed analytical tools that include Disciple-LTA (Tecuci et al., 2008; Schum et al., 2009), TIACRITIS (Tecuci et al., 2011), Disciple-CD (Tecuci et al., 2016a) and Cogent (Tecuci et al., 2015; 2018).

Deep Anticipatory Intelligence Analysis in the Framework of the Scientific Method

An expert analyst directly teaches the instructable agent how to perform anticipatory analysis in a way that is similar to teaching a student. The analyst follows the systematic analysis process summarized in Figure 1 that is inspired by the *scientific method of hypothesis generation and testing*. This is a mixed-initiative process (Tecuci et al., 2007b) where the analyst and the instructable agent: (1) Use *abductive reasoning* that shows that something is *possibly* true to generate alternative anticipatory hypotheses that may explain an *alert* (an indicator of a situation of interest), or are the alternative answers of an *intelligence question*; (2) Use each hypothesis and *deductive reasoning* that shows that something is *necessarily* true to guide the *discovery and collection* of additional evidence; and (3) Use *inductive reasoning* that shows that something is *probably* true to assess the *probabilities* of the hypotheses and the *confidence* in these probabilities based on the collected evidence. As a result, the agent learns general rules to generate hypotheses, rules to search for evidence, and rules to assess the hypotheses that enable it to reason autonomously.

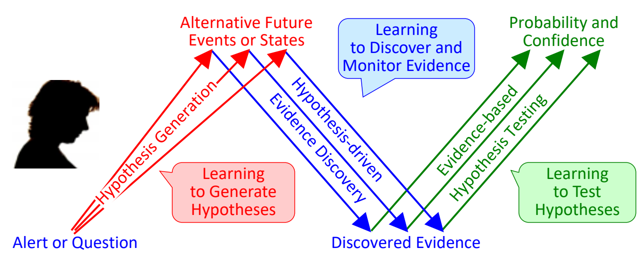

Figure 1. Model of the analysis process.

Hypothesis Generation

Figure 2 is an abstract representation of deep anticipatory analysis which is a multi-step iteration of the process from Figure 1. The generation of each hypothesis of interest (e.g., H_o and H_p at the top of Figure 2) is done through a chain of abduction steps (i.e., $Alert \rightarrow F_i \rightarrow G_l \rightarrow H_o$ and $Alert \rightarrow F_i \rightarrow G_n \rightarrow H_p$) where each hypothesis in the chain is either at least barely likely (50-55%) or high risk (Peirce, 1955; Schum, 2001a). Notice that this process is much more efficient than that based on the single-step abduction $Alert \rightarrow H_o$ that would require the analysis of H_o and of all its many alternatives represented by the dots at the top of Figure 2. We call it *deep* anticipatory analysis because it involves the analysis of all the hypotheses in Figure 2.

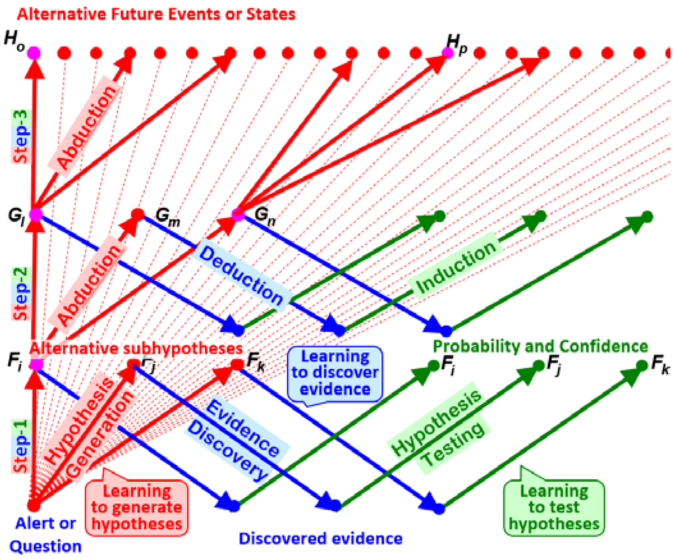

Figure 2: Deep anticipatory analysis.

Evidence Discovery

During the second phase, *hypothesis-driven evidence discovery*, one decomposes each of the generated hypotheses to discover new evidence, as illustrated with hypothesis H in Figure 3. The guiding question is: What evidence would favor or disfavor H? To answer it, one develops a Wigmorean probabilistic inference network (Wigmore 1937; Schum, 2001b; Tecuci et al., 2016a), successively decomposing H into simpler and simpler hypotheses by considering probabilistic favoring (under the left, green square) and disfavoring arguments (under the right, pink square). For example, $H_1\&H_2$ is a probabilistic favoring argument for H (if H_1 is true and H_2 is true, then it is likely that H is true). H_3 is a probabilistic disfavoring argument (if H_3 is true, then it is very likely that H is false). Similarly, H_{1a} and H_{1b} are two probabilistic favoring arguments for H_1. Such successive decompositions, through favoring and disfavoring arguments, continue until the resulting leaf hypotheses are simple enough to point to what evidence would favor or disfavor them (e.g., ***Collect*** ***evidence*** *to determine whether* H_{1a}). We use a refinement of the ICD 203 (2015) probability scale where we have split ICD-likely (55-80%) into *likely* (55-70%) and *more than likely* (70-80%), in order to enable more precise assessments.

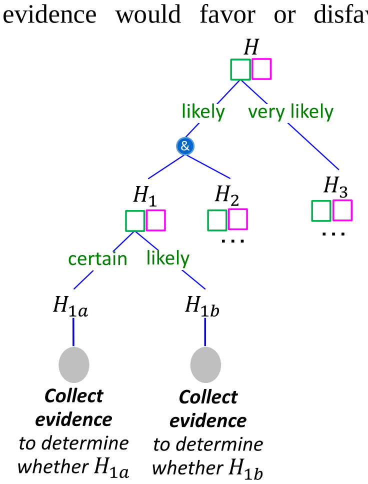

Probability Scale	
100%	certain
95-99%	almost certain
80-95%	very likely
70-80%	more than likely
55-70%	likely
50-55%	barely likely
45-50%	barely unlikely
30-45%	unlikely
20-30%	less than unlikely
5-20%	very unlikely
1-5%	almost no chance
0%	no chance

Figure 3. Evidence discovery.

Hypothesis Testing

Finally, one fuses the discovered evidence to assess the *probability* of each hypothesis and the *confidence* in this probability. First, one assesses the probabilities of the bottom-level hypotheses and the confidence in these probabilities based on the discovered evidence, as illustrated in Figure 4. In this case there is a single item of evidence E favoring the hypothesis H. One can assess the probability of hypothesis H based on E by assessing the three credentials of evidence -- *credibility*, *relevance*, and *inferential force*. The *credibility* of evidence answers the question: *What is the probability that the evidence is true?* The *relevance* of evidence answers the question: *What would be the probability of the hypothesis if the evidence were true?* The *inferential force* of evidence answers the question: *What would be the probability of the hypothesis based only on this evidence?* Each probability assessment is paired with an assessment of the confidence in that probability. Table 1 presents the definition of the confidence levels, which are based on the definitions from (Joint Chiefs of Staff, 2013; DIA, 2015; NIC, 2017) where we introduced two additional values, *very low confidence* and *very high confidence*. Then, the probability and confidence of each upper-level hypothesis is computed based on the logical structure of the argumentation (conjunctions and disjunctions of hypotheses), using the *min-max* probability combination rules common to the *Baconian probability view* (Cohen, 1977) and the *Fuzzy view* (Zadeh, 1983). *On-balance judgements* are made for the favoring and disfavoring arguments of a hypothesis. The min-max rules are simpler and more intuitive than the *Bayes rule* used in the *Bayesian probabilities view*, or the *Dempster-Shafer rule* used in the *Belief Functions view* (Tecuci et al., 2016a, pp.177-196).

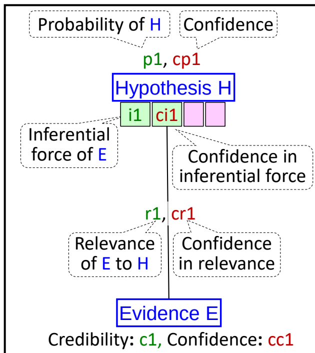

Figure 4: Evidence credentials.

From the literature on intelligence analysis, as well as the intelligence and defense organizations posited guidelines for expressing uncertainty (DIA, 2015; JCS, 2013; NIC, 2017), we have identified the following criteria for the confidence in a probabilistic assessment: *(1) The reliability/credibility of the sources used* -- acknowledged by everyone; *(2) The degree of corroboration of the evidence* -- widely acknowledged; (3) *The number of assumptions made and their influence on the analytic conclusion* (DIA, JCS); (4) *The intelligence gaps remaining* (DIA, JCS, NIC); (5)

Potential for deception (DIA); (6) *Responsiveness to new information* -- the less the analysts expect subsequent evidence and analysis to change their judgments the more confident they are in the current analysis and (7) *Range of reasonable opinion* -- the narrower the set of plausible viewpoints analysts have the more confident they are in their assessments (Friedman and Zeckhauser, 2015); (8) *Time pressure and task complexity* and (9) *Level of analyst collaboration* (Peterson, 2008); (10) *Reasoning / Strength of analytic inferences / Quality of logical inferences and analytic methods used* (Peterson; DIA, JCS, NIC).

We are investigating a method to assess the confidence in probability by taking into account all these criteria.

Table 1. Confidence levels.

Total confidence: Known true reasoning.

Very High: Very high quality and very well corroborated information from proven sources, no potential for deception, no assumptions or gaps, sound reasoning.

High: High quality and well corroborated information from proven sources, low potential for deception, non-critical assumptions and/or gaps, undisputed reasoning.

Moderate: Partially corroborated information from good sources, moderate potential for deception, potentially critical assumptions used to fill gaps, or a mix of inferences.

Low: Uncorroborated information from good or marginal sources, high potential for deception, key assumptions used to fill critical gaps, or mostly weak inferences.

Very Low: Uncorroborated information from marginal sources, very high potential for deception, key assumptions used to fill critical gaps, and weak inferences.

No confidence: Known false reasoning.

Agent Architecture

Figure 5 shows the envisioned architecture of the instructable agent for anticipatory analysis. The analyst demonstrates to the *Mixed-Initiative Learning and Reasoning Assistant* how to perform a specific anticipatory analysis by following the process in Figure 1. Then an ontology of the entities referred in the analysis is defined (Allemang et al., 2020; W3C, 2004). After that the agent learns reasoning rules as ontology-based generalizations of the individual reasoning steps from the analysis, as illustrated in Figure 6. Ontology development and rule learning are based on the Disciple multistrategy apprenticeship learning approach (Boicu et al., 2001;

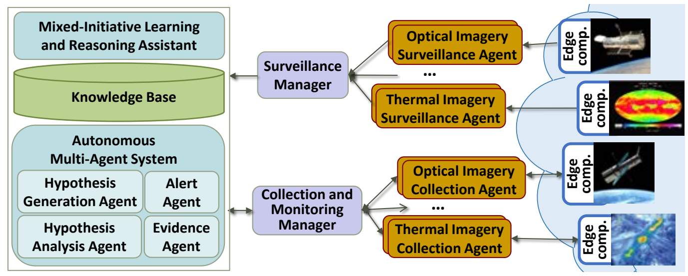

Figure 5. Agent architecture.

Tecuci, 1988; 1998; Tecuci and Hieb, 1996; Tecuci et al., 2000; 2002; 2005; 2007a; 2016b; 2019; Huang et al., 2020).

Notice that the reasoning step in the upper-left side of Figure 6 is an illustration of the leaf steps in Figure 3. The ontology fragment is an ontology of collection agents and their functions.

A *collection agent* uses a *convolutional neural network (CNN)* that needs to be *trained* to perform its function. For example, with input from an *optical imagery sensor* a CNN can detect whether there are plumes of smoke over an entity plant (Ba et al., 2019). The output from the collection agent will be an item of evidence detected with an accuracy characteristic to the trained CNN, for example, 92.75% (i.e., very likely):

> E1 Plumes of smoke (optical imagery sensor detected plumes of smoke over the Tanan plant as of 6/15/2021) with accuracy very likely (92.755%) and confidence very high.

The ontology and the learned rules are stored in the Knowledge Base which is shared with the Autonomous Multi-Agent System, enabling it to automatically respond to alerts by performing anticipatory analysis and presenting the conclusions to an *on-the-loop* analyst.

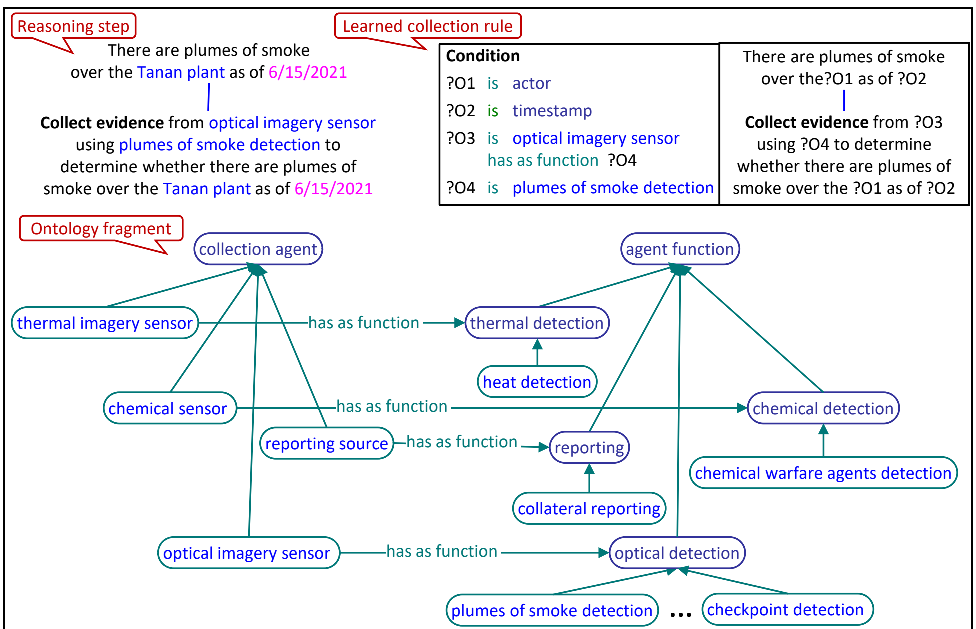

Figure 6. Ontology-based rule learning.

References

Allemang, D.; Hendler, J.; and Gandon, F. 2020. *Semantic Web for the Working Ontologist: Effective Modeling for Linked Data, RDFS and OWL*. Morgan Kaufmann Publishers.

Ba, R.; Chen, C.; Yuan, J.; Song, W.; and Lo, S. 2019. SmokeNet: Satellite Smoke Scene Detection Using Convolutional Neural Network with Spatial and Channel-Wise Attention. *Remote Sensing*, 11(4), 1702. doi.org/10.3390/rs11141702

Boicu, M.; Tecuci, G.; Stanescu, B.; Marcu, D.; and Cascaval, C.E. 2001. Automatic Knowledge Acquisition from Subject Matter Experts. In P*roc. of the 13th Int. Conf. on Tools with Artificial Intelligence (ICTAI)*, November 7-9, Dallas, TX: 69-78.

Cohen, L.J. 1977. *The Probable and the Provable.* Oxford: Clarendon Press.

DIA. 2015. *Defense Intelligence Agency Guidelines*.

Friedman, J.A., and Zeckhauser, R. 2015. Analytic Confidence and Political Decision Making: Theoretical Principles and Experimental Evidence from National Security Professionals. *Political Psychology* 39 (5). DOI: 10.1111/pops.12465

Heuer, R.J., and Pherson, R.H. 2011. *Structured Analytic Techniques for Intelligence Analysis*. Washington, DC: CQ Press.

Huang, J.; An, Z.; Meckl, S.; Tecuci, G.; and Marcu, D. 2020. Complementary Approaches to Instructable Agents for Advanced Persistent Threats Detection. *Studies in Informatics and Control*, 29 (3).

ICD 203. 2015. Analytic Standards. *Intelligence Community Directive 203*. fas.org/irp/dni/icd/icd-203.pdf

Joint Chiefs of Staff. 2013. *Joint publication 2–0: Joint intelligence*. Washington, DC.

Klein, G.; Snowden, D.; and Pin, C.L. 2007. Anticipatory Thinking. In K. Mosier, K., and U. Fischer, U. (eds), *Proceedings of the Eighth International NDM Conference*. Pacific Grove, CA.

Marrin, S. 2011. *Improving Intelligence Analysis: Bridging the gap between scholarship and practice*. London & New York: Routlege.

NIC. 2017. *National Intelligence Council Guidelines*.

ODNI. 2019. *National Intelligence Strategy of the United States of America*. Office of the Director of National Intelligence.

Peirce, C.S. 1901. Abduction and Induction. In Buchler, J. (editor) 1955. *Philosophical Writings of* Peirce: 150-156. New York, NY: Dover Publications.

Peterson, J. 2008. Appropriate Factors to Consider when Assessing Analytic Confidence in Intelligence Analysis. *Master of Science in Applied Intelligence*, Mercyhurst College.

Pherson, K.H., and Pherson, R.H. 2021. *Critical Thinking for Strategic Intelligence.* SAGE Publishing.

Schum, D.A. 2001a. Species of Abductive Reasoning in Fact Investigation in Law. *Cardozo Law Review*, 22 (5-6): 1645-1681.

Schum, D.A. 2001b. *The Evidential Foundations of Probabilistic Reasoning*. Northwestern University Press.

Schum, D.; Tecuci, G.; and Boicu, M. 2009. Analyzing Evidence and Its Chain of Custody: A Mixed-Initiative Computational Approach. *International Journal of Intelligence and CounterIntelligence*, 22 (2): 298-319.

Tecuci, G. 1988. DISCIPLE: A Theory, Methodology and System for Learning Expert Knowledge. *These de Docteur en Science* (in English), 197 pages, University of Paris-South, France.

Tecuci, G. 1998. *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*. San Diego, CA: Academic Press.

Tecuci, G., and Hieb, M.H. 1996. Teaching Intelligent Agents: The Disciple Approach. *International Journal of Human-Computer Interaction*, 8 (3): 259-285.

Tecuci, G.; Boicu, M.; Bowman, M.; Marcu, D.; Shyr, P.; and Cascaval, C. 2000. An Experiment in Agent Teaching by Subject Matter Experts. *International Journal of Human-Computer Studies*, 53: 583-610.

Tecuci, G.; Boicu, M.; Marcu, D.; Stanescu, B.; Boicu, C.; Comello, J.; Lopez, A.; Donlon, J.; and Cleckner, W. 2002. Development and Deployment of a Disciple Agent for Center of Gravity Analysis. In *Proc. of the 18th National Conf. of Artificial Intelligence and the 14th Conf. on Innovative Applications of Artificial Intelligence*, Edmonton, Alberta, Canada: 853-860.

Tecuci, G.; Boicu, M.; Boicu, C.; Marcu, D.; Stanescu, B.; and Barbulescu, M. 2005. The Disciple-RKF Learning and Reasoning Agent. *Computational Intelligence,* 21: 462-479.

Tecuci, G.; Boicu, M.; Marcu, D.; Boicu, C.; Barbulescu, M.; Ayers, C.; and Cammons, D. 2007a. Cognitive Assistants for Analysts. *Journal of Intelligence Community Research and Development (JICRD)*. Also in Auger, J., and Wimbish, W. (editors), *Proteus Futures Digest: A Compilation of Selected Works Derived from the 2006 Proteus Workshop*, 303-329.

Tecuci, G.; Boicu, M.; and Cox, M.T. 2007b. Seven Aspects of Mixed-Initiative Reasoning: An Introduction to the Special Issue on Mixed-Initiative Assistants. *AI Magazine*, 28 (2): 11-18.

Tecuci, G.; Boicu, M.; Marcu, D.; Boicu, C.; and Barbulescu, M. 2008. Disciple-LTA: Learning, Tutoring and Analytic Assistance. *Journal of Intelligence Community Research and Development*.

Tecuci, G.; Marcu, D.; Boicu, M.; Schum, D.A.; and Russell, K. 2011. Computational Theory and Cognitive Assistant for Intelligence Analysis. In *Proceedings of the 6th International Conference on Semantic Technologies for Intelligence, Defense, and Security - STIDS 2011*, November 16-18, Fairfax, VA: 68-75.

Tecuci, G.; Marcu, D.; Boicu, M.; and Schum, D.A. 2015. Cogent: Cognitive Agent for Cogent Analysis. In *Proceedings of the 2015 AAAI Fall Symposium "Cognitive Assistance in Government and Public Sector Applications,"* Arlington, VA. Technical Report FS-15-02: 58-65, Palo Alto, CA: AAAI Press.

Tecuci, G.; Schum, D.A.; Marcu, D.; and Boicu, M. 2016a. *Intelligence Analysis as Discovery of Evidence, Hypotheses, and Arguments: Connecting the Dots*. Cambridge University Press.

Tecuci, G.; Marcu, D.; Boicu, M.; and Schum, D.A. 2016b. *Knowledge Engineering: Building Cognitive Assistants for Evidence-based Reasoning*. Cambridge University Press.

Tecuci, G.; Kaiser, L.; Marcu, D.; Uttamsingh, C.; and Boicu, M. 2018. Evidence-based Reasoning in Intelligence Analysis: Structured Methodology and System. *Computing in Science and Engineering,* 20 (6): 9-21.

Tecuci, G.; Meckl, S.; Marcu, D.; and Boicu, M. 2019. Instructable Cognitive Agents for Autonomous Evidence-Based Reasoning. *Advances in Cognitive Systems,* 8.

Wigmore, J. H. 1937. *The Science of Judicial Proof: As Given by Logic, Psychology, and General Experience and Illustrated in Judicial Trials, 3rd edition*. Boston, MA: Little, Brown & Co.

W3C. 2004. www.w3.org/TR/rdf-schema/

Zadeh, L. 1983. The Role of Fuzzy Logic in the Management of Uncertainty in Expert Systems. *Fuzzy Sets and Systems*, 11: 199-227.