A Conversational Intelligent Agent for Career Guidance and Counseling

Andrew J. Hampton,[1] Vasile Rus,[1] Frank Andrasik,[1] Benjamin D. Nye,[2] Arthur C. Graesser[1]

[1]University of Memphis, Institute for Intelligent Systems

[2]University of Southern California, Institute for Creative Technologies
andrew.hampton@memphis.edu

Abstract

We present the preliminary construction of a mixed initiative conversational intelligent agent designed to provide guidance on career pathway information and available resources for common career-relevant personal problems. We use a single job classification within the United State Navy as a proof of concept. With this system, user input from career guidance sessions is linked via advanced natural language processing techniques to our framework of Navy training and education standards, promotion protocols, and organizational structure, producing feedback on resources and recommendations sensitive to user history and stated career goals. In recognition of the variety of personal problems that can impact career progress, the intelligent agent also offers rudimentary "counseling". Detection of language related to issues (e.g., sleep deprivation, financial difficulties, substance abuse) triggers targeted dialogues that gather more information, offer tailored suggestions, and/or provide referrals to appropriate resources such as a human counselor when in-depth counseling is warranted. Sessions occur when sailors initiate them, when performance or progress drops below Navy expectations, or with respect to career milestones (the latter two enabled by integration with career monitoring databases). This software, currently in alpha testing, has the potential to serve as an intuitively accessible information hub, engaging and encouraging sailors to take ownership of their career paths in the most efficient way possible, benefiting both individuals and the Navy as a whole.

Introduction

Navigating a career constitutes one of life's most enduring challenges, particularly within a massive organization like the United States Navy. Though the Navy has numerous resources for guidance, accessing and identifying key information sources across the many existing platforms can be challenging for sailors (e.g., determining the appropriate program or point of contact, developing an accurate understanding of processes, and even recognizing the need for planning itself). Focusing on evaluations, education, certifications, and training is quite demanding, even before considering their cumulative long-term implications. Further, potential for success can be restricted by generic personal issues, such as financial difficulties, or Navy-specific issues, like homesickness when at sea.

The dizzying breadth of factors that inform career success, combined with the glut of disparately formatted databases and websites, suggests the need for a central hub of information. This hub should enable exploration of resources for career and relevant personal challenges by first discovering what factors are most important to the sailor. As far as possible, this interface should be friendly and welcoming to counteract the vast array of impersonal static information. It should also remember user input and preferences to provide a touchstone over long stretches without use, reminding users of stated goals and progress relative to them. This kind of friendly, relatively open-ended career tool would then be ideally positioned to accept input of a more personal nature which may have direct bearing on career prospects

We present in this paper the theoretical foundations, design, and promise of a mixed-initiative, conversational intelligent system for career guidance and counseling—*Claire the Counselor*. This conversational intelligent agent has been designed to provide a user-friendly, interactive environment that recognizes user input pertinent to Navy career navigation issues and provides guidance to appropriate resources within the Navy. We link user input from "counseling sessions," leveraging advanced natural language processing techniques, to our framework of Navy training and education standards, career advancement protocols, and organizational structure, producing feedback on resources and recommendations sensitive to user history and stated career goals. An instance of the proposed system has been prototyped and will be illustrated for career guidance in one rate (i.e., job classification) of the U.S. Navy, namely avionics electronics technician. This is meant to constrain the possible career states, transitions, and trajectories to a

manageable subset, while enabling explorations of issues common to all rates. Though the infrastructure underlying our design currently remains restricted, in this manuscript we will primarily refer to generic career characteristics.

The system is mixed-initiative, meaning both the system and the user can initiate a counseling session. Integration with an existing life-long learning program (PAL3; Swartout et al. 2016) enables system-initiated sessions. Within that broader life-long learning software, the capacity to monitor sailors' career progress affords proactively triggering sessions at key moments, such as before major career milestones or when performance drops below Navy expectations. Other types of system-triggered sessions involve positive feedback and informative dialogues, e.g., offering compensation related information (drawing upon existing Navy career guidance protocols). Together, these address the problem of young professionals not necessarily knowing when critical career inflection points are occurring. Sailors may also wish to proactively check on or adjust their career trajectory. For instance, sailors may initiate a session to learn what they need to do to move to a leadership position.

In recognition of the role personal difficulties play in career progression, and in the similarly distributed set of resources available for mediation, our intelligent agent also offers counseling for personal problems. Upon recognizing language pertinent to a pre-defined set of personal issues, the system initiates targeted dialogues designed to gather more information, offer tailored suggestions, and provide referrals to appropriate resources, including human counselors when in-depth counseling is warranted. Users are advised that all counseling sessions are strictly confidential, which is designed to increase trust and encourage use of the proposed system. Even updates to the Learner Model are by user consent only. The current prototype, in alpha testing, has the potential to serve as a centralized information hub, engaging and encouraging sailors to take ownership of their career paths in the most efficient way possible, benefiting both individuals and the Navy as a whole.

Next, we present an overview of the design of an intelligent counseling agent. Then, we describe in more depth our solutions to core problems when developing dialogue-based counseling systems: speech act classification and understanding user responses. We end the paper with conclusions and future work.

An Intelligent Counseling Agent

Our goal has been to develop a virtual agent software system which provides career guidance and counseling services through a personalized, mixed-initiative, interactive, and engaging process. The general approach leverages our experience with building virtual agent-based training systems, making use of discourse and dialogue theories and

principles that our group has researched for decades (Rus, D'Mello, Hu, and Graesser 2013), while building upon Navy resources and lessons learned from previous work such as the 5 Vector Model (NPDC 2005), an ambitious but ultimately failed attempt to classify characteristics of sailors (skills, abilities, competencies, certifications, etc.) relevant to job assignments.

More concretely, we were guided by the following specific goals:
• Offer persistent guidance and mentoring across career transitions
• Enable continuity across and enable smooth career and assignment transitions
• Build a long-term relationship with each sailor by establishing rapport and trust
• Provide anytime/anywhere counseling and guidance service accessible from any Internet-connected device
• Use electronics training and on-the-job experiences data to guide the counseling dialogues
• Recommend existing high-quality learning resources
• Offer counseling on common "life problems" through a motivational interviewing framework

The Learner Model component captures sailors' knowledge state (knowledge, skills, abilities, and competencies) and other characteristics such as personality traits. The knowledge state is inferred from various sources such as report cards from school (K–12), the Armed Services Vocational Aptitude Battery (ASVAB), on-the-job experiences and performance, and annual evaluations. Inclusion (and form) of these depends upon Navy recording, distinct from Claire's core functionality. Sailors' wishes and interests are also considered when offering career guidance and counseling as well as specific job requirements that are available for every Navy rating. This influence is critical in providing a tool that sailors feel creates personal benefit, particularly with the Navy's focus on retention beyond initial tours of duty.

Consideration for the Navy's needs ensures pragmatic constraints on recommendations with respect to allocation of training and personnel resources. This can help optimize usage and reduce turnover. Navy needs can be specified in terms of job openings for each career, e.g. the number of cooks needed at some time. Requirements for a specific job are specified as sets of knowledge, skills, abilities, and competencies needed to perform the duties related to the job. The design of the counseling system was driven by the need to optimize the match between sailors' capabilities (skills, abilities, competencies, knowledge), wishes, and interests, on one hand, and Navy needs on the other. Critically, this job assignment optimization is done globally, across all sailors and Navy needs. When system-wide matching optimization is achieved, Navy mission readiness will in turn be maximized.

The job assignment optimization problem is one of the fundamental combinatorial optimization problems and consists of finding a maximum weight matching in a weighted bipartite graph. Given a complete bipartite graph, G = (S, T, E), with n sailor vertices (S), n ships vertices (T), and each edge $e_{s \in S, t \in T} \in E$ has a non-negative weight $w(s, t)$ indicating how qualified a sailor is for a certain job, the task is to find a matching M from S to T with maximum weight. In case of different numbers of sailors or ships, dummy vertices could be used.

The assignment problem can be thus formulated as finding a permutation π for which $SOPT = \sum n_{i=1} w(s_i, t_{\pi(i)})$ is maximum. Such an assignment is called optimum assignment. An algorithm, the Kuhn Munkres method (Kuhn 1955), has been proposed that can find a solution in polynomial time (see Dawes 2011, for a complete formal description of the algorithm).

Worth mentioning is the fact that a best matched assignment does not mean a perfect match. For example, although one sailor's profile might best match the requirements of an assignment among all candidate sailors, that sailor might lack key job requirements or need re-training due to, for instance, skill decay in situations when the sailor has not applied the skill for an extended period of time. (Modeling and mitigation of skill decay is more directly addressed by the PAL3 program with which Claire is integrated. See Hampton et al. 2018 for details.) The counseling agent can provide specific guidance to help the sailor be ready for the job. In most cases, the agent can recommend learning resources to address any skill/knowledge/competencies gaps or decays therein.

Many other scenarios exist in which a counseling agent could be extremely useful. Currently, Claire the Counselor executes the following counseling scenarios: (1) information providing dialogues with respect to general career path for specific ratings; (2) precise information providing dialogues for frequently asked questions such as the salary associated with a given level for a particular; (3) counseling for someone who is assigned or would like to serve in a leadership role; and (4) counseling on common "life" problems through a motivational interviewing framework.

Tasks (1), (2), and (3) are implemented as informative dialogues and require access to Navy career information in machine readable form. This is challenging because career information is complex, distributed, dynamic, diverse, and spans many decades. A capable system must have information extraction components that can retrieve relevant documents from disparate Navy sources (described later) and then extract from those documents information relevant to specific career-related questions, anticipating sailor inquiries. Details of this process are presented next.

Near-, Medium-, and Long-Term Career Guidance

The overarching goal is for an artificially intelligent career guidance agent to guide sailors throughout their entire careers in the Navy by helping them set learning objectives based on Navy needs, (e.g., job requirements for their next assignment), and their current knowledge, skills, abilities, and competencies as reflected in the Learner Model. To this end, we designed our system to offer near-, medium-, and long-term career guidance. Again, all counseling sessions are strictly confidential. When combined with the reduced inhibition inherent in mediated communication (Kiesler, Siegel, and McGuire 1984), this feature can lead to increased trust and use of a counseling system.

An example of a near-term career goal would be when a sailor is given a new assignment on a new platform, e.g., electrical systems pertaining to the FA-18 jet starting in three months, for which they need to update their knowledge and skills on how to diagnose and repair that machine's avionics. The counselor agent would recommend specific learning resources that the sailor could use. By the same token, the system could monitor electronic maintenance tracking data and on the job experiences to make recommendations. For instance, if someone is supposed to perform radar repairs but has struggled with diagnosing that equipment in the recent past, the system should suggest remedial resources tailored to the sailor's level of proficiency, providing a constantly engaging but never overwhelming learning resource. Training and assessment, like all content-based resources in Claire, would be handled by an external resource and subsequently update the Learner Model.

A medium-term career planning example would be a situation in which someone qualifies for promotion in a year but needs leadership credentials to improve their position among peers. In this case, a sufficiently adept counseling agent would recommend leadership learning resources. Long-term career guidance can take the form of presenting the major career milestones and average time to reach those milestones for a Navy rating such as aviation electronics technician. It takes 22 years on average to reach master chief for this rating and knowing when and how to prepare for the next milestone is of great importance.

The system initiative can spawn from information being pushed on it, e.g., when a significant sailor achievement such as earning a cybersecurity certification is recorded in the Learner Model. Alternatively, information can be pulled from the Learner Model or from an alternative distinct career goals and training objectives database that we intend to build (possibly using a commercial Learning Record Store, LRS, such as Learning Locker). A push-and-pull protocol between the LRS and a career counseling system is needed to provide timely and effective counseling.

Significant potential for optimization readily suggests itself, such as data mining of ideal learner profiles (i.e., where an exceptional sailor should be in his or her career development) and learner trajectories that combine to provide optimal counseling. The technical and practical challenges of these improvements are commensurate with their promise, and are therefore beyond the scope of this manuscript. However, some of those additional components are under development as of this writing.

Frequently Asked Questions

To supplement complex knowledge representation and interactive recommendations, we compiled a list of frequent queries regarding general Navy life from several popular Navy support websites. These include both official, affiliated, and unaffiliated (but credible) resources that cover a wide range of information from how long an average deployment lasts to the availability of financial aid for service members. We eliminated duplicate questions and any that could not be answered in a relatively brief response to produce 31 questions and answers broken into sections including "Becoming an Officer", "The Navy Reserves", and "Navy Life". Sailors can access this list of questions at any time via an "FAQ" link on the interface. A conversational version of this interface is under construction. These "canned" responses augment Claire's capacity to respond to free sailor input and increase her perceived intelligence.

The creation of a more complete FAQ library will naturally follow extended use of the system by the target population. As they ask certain questions frequently, we can identify these and produce appropriate answers with trivial effort. The FAQs are triggered by either automatic detection of such questions during conversations, in which case the system simply provides the answer associated with that frequently asked question, or through an on-screen button that users can click at any moment and browse the entire list of FAQs.

Motivational Interviewing

Our system also recognizes language indicative of common "life" problems, not specific to the Navy, such as financial troubles, depression, or substance abuse. The Navy typically has exceptional support for these issues, but sailors may not know where to find and access them. An anonymous conversational interface designed to detect either tangential or pointed reference to these difficulties can lower the barrier to self-awareness and self-care, allowing sailors to explore resources completely free of judgment or obligation. Our system can refer the sailor to the appropriate Navy resources that address problems of immediate concern.

Motivational Interviewing (MI) (e.g., Miller and Rollnick 2012) serves as the guiding framework for counseling on common "life" problems. MI arose as a means of addressing several problems individuals face when, on one hand they are aware that a change in behavior is needed, yet, at the same time lack in motivation and information, or else and experiencing denial and ambivalence about changing. In these types of situations, approaches that are too directive, confrontational, or appeal to expert advice fall upon deaf ears. MI is best characterized as non-adversarial, nonconfrontational, and non-judgmental. Rather, it seeks to engage individuals in becoming more aware of their conflicts, resultant problem behaviors, and the consequences that ensue if one "stays the course." MI encourages them to begin to consider more appropriate choices while enhancing motivation for change.

While most investigations of MI have centered upon lifestyle modification and chronic diseases, Stoltz and Young (2012) have adapted this approach to career counseling. Our focus on counseling active duty personnel is guided and informed by their work, plus the extensive literature base that now exists on MI (a recent search of PubMed identified approximately 3,500 citations).

Stoltz and Young distinguish seven distinct stages for applying MI in career counseling: (1) building a working (therapeutic) relationship, (2) assessing readiness to change, (3) identifying and discussing discrepancies, (4) "rolling with resistance" as appropriate (refraining from disagreeing, exploring sources of ambivalence), (5) promoting "self-efficacy", (6) expressing empathy (drawing upon Rogerian or person-centered techniques of accurate reflection) (e.g., Roger 1959), and (7) preparing for transition (and termination). When considering all that is involved, one might assume the process to be long and involved.

However, notable outcomes have been achieved with varied target problems, ranging from a single, brief session to five or more (Rubak et al. 2005). We have implemented in Claire the Counselor a streamlined version of this seven-stage process to maintain a tenable number of encounters. To this end, our MI sessions include generic question asking, reflective listening statements, tactful advice giving and feedback sharing, and key summarizing statements that help inform our career counseling approach. The dialogue manager has been designed to include such dialogue moves and policies on when to trigger them, as described in the next section.

Dialogue Management and Understanding User Input

Claire the Counselor's main interaction with the user is dialogue-based, mimicking the interaction between a sailor and a human career guidance and counseling expert. Once the

user logs into the system, the system greets and tries to quickly identify the user's key issue. Right from the start, two major technical challenges must be addressed to automatically manage the conversation-based interaction with the user: identifying user's intentions and understanding their natural language responses.

Speech Act Classification

An important component of the dialogue manager is the identification of user's intentions based on their utterances; i.e., the task of speech act classification (SAC; Rus et al. 2012, 2013). The SAC uses a multi-leveled taxonomy of speech act categories with 4 major categories at the top (meta-cognitive, meta-communicative, question, and contribution) and 35 categories total at the second and third level. When a user utterance is labeled as being a contribution, which is a content-rich statement, it is passed on to the semantic processing component for a deeper understanding. Contributions from users can be relevant or irrelevant (e.g., a content-rich statement that talks about food or friends when the topic is Navy careers).

Only relevant contributions are passed on for a deeper understanding. Some categories of student responses, e.g., Greetings, can be handled immediately without the need for deep semantic understanding. For instance, to a Greeting from the user, the system can automatically respond with a typical greeting expression, if it has not already greeted the user, or can respond with a dialogue move that pushes the dialogue forward, e.g., asking "How may I help you?" at the beginning of a session in order to elicit from the user responses that indicate the main issue they need help with. When the system initiates a session, the system already recognizes the issue needing attention and therefore could start by presenting the issue to the user. These dialogue policies are implemented as of this writing using a state-transition network and production rules. However, as experimental data become available from actual users interacting with Claire the Counselor, we hope to be able to infer dialogue policies automatically through data mining and machine learning methods, such as reinforcement learning.

Understanding Users' Natural Language Responses

Algorithms are needed to interpret the meaning of students' natural language contributions at each turn in the dialogue. This section describes advances including the addition of negation handling and syntactic information as well as proposing algorithms that incorporate optimized semantic matching solutions.

Semantic similarity is the underlying principle for understanding users' contributions. That is, we understand a user response by assessing how similar it is to an expert answer. The meaning of the expert answer is deemed known and therefore the student contribution is deemed known if it is semantically similar to the expert answer. The alternative to the semantic similarity approach is the true or full understanding approach in which the student response is fully interpreted. The full understanding approach is intractable for real-world, sizeable applications, such as Claire the Counselor, because it requires vast amounts of world and domain knowledge. Domain and world knowledge are captured to some extent by semantic similarity approaches.

Because user contributions can vary in length, from minimal, (e.g., including just one content word such as "equal"), to a sentence or even a paragraph, we need methods that apply across different granularities of texts. The methods presented next are generally applicable to texts of various sizes although some are more suited for a certain granularity level. For instance, when syntactic information is used in a particular method then the method cannot be applied at word level directly. Several broad categories of semantic similarity methods were investigated and are included in our SEMILAR toolkit: vectorial methods including LSA (Landauer et al. 2007; Lintean, Moldovan, Rus, and McNamara 2010), probabilistic methods including Latent Dirichlet Allocation (LDA; Blei, Ng, and Jordan 2003; Niraula et al. 2013), greedy methods, optimal methods (Rus et al. 2012; Rus and Lintean 2012), and some others. Due to space reasons, we are not able to present more details about all these methods. We use a variant of our optimal word-to-word and syntactic matching via Quadratic Assignment (QAP; Koopmans and Beckmann 1957; Lintean and Rus 2015). This method finds an optimal global assignment of words in one sentence (e.g., a user response) to words in the other sentence (e.g., the expert answer) based on their word-to-word similarity, while simultaneously maximizing the match between the syntactic dependencies between words.

A brute force solution to the QAP problem, which would generate all possible mappings from words in a sentence to words in another sentence, is infeasible because the solution space is too large. For example, when considering all possible pairings of words between sentence A, of size n, and sentence B of size m, where n < m, and we pose no limitations on the type of pairings that can be made, $m!/(m - n)!$ possible solutions exist. For sentences of average size $n = m$ = 20 words, there are 2.4 * 1018 possible pairings.

An efficient branch-and-bound algorithm has been developed to reduce the explored space in search of the optimal solution. This is possible by defining a bounding function that always overestimates or underestimates solutions, depending on what type of optimal solution is sought, maximum or minimum cost, respectively.

Conclusions and Next Steps

We presented the design of a career guidance and counseling agent that can help sailors navigate career paths in the Navy as well as recommend learning resources to keep their knowledge, skills, and competencies updated and aligned with their job assignments. Counseling on common life problems supplements these in recognition of their impact on careers and the similar structure of disparate resources that sailors may not know they need or have access to. Given the distributed, dynamic, and complex nature of Navy careers, there is an acute need for automated career guidance and counseling, such as Claire the Counselor, that could be extremely useful to sailors who have limited time for information retrieval while performing their already demanding duties. Conversation-based interaction and availability 24/7 will make access to career-related information much more accessible, increasing sailor job satisfaction and readiness by providing recommendations for learning resources in a timely and well-planned manner. Integration with the more near-term focused PAL3 software ensures at least some level of regular exposure to Claire. Our plans include testing the system through several experiments and refining aspects as needed based on the outcomes of those experiments before deploying it on a large scale in the Navy.

It should be noted that while our current focus was on only one Navy rating, Avionics Electronics Technician, the design is flexible and scalable to all Navy ratings, which we plan to pursue in the future. Developing a career guidance and counseling agent is a major endeavor and therefore a full and detailed specification of the design is not possible in a limited-space paper such as this one. We focused instead on the main features and some of the major algorithms. Further, this proof-of-concept would be readily amenable to application in other broad but finite fields with predominantly young professionals navigating within broadly defined categories, e.g., majors in a university.

Acknowledgments

References

Blei, D. M., Ng, A. Y., and Jordan, M.I. 2003. Latent Dirichlet Allocation. *The Journal of Machine Learning Research* 3, 993–1022.

Hampton, A. J., Nye, B. D., Pavlik, P. I., Swartout, W. R., Graesser, A. C., and Gunderson, J. 2018. Mitigating Knowledge Decay from Instruction with Voluntary Use of an Adaptive Learning System. In *Proceedings of the 2018 Artificial Intelligence in Education Conference*, 119–133. London, England: Artificial Intelligence in Education.

Kiesler, S., Siegel, J., and McGuire, T. W. 1984. Social Psychological Aspects of Computer-mediated Communication. *American Psychologist* 39(10): 1123–1134.

Koopmans, T. C., and Beckmann, M. 1957. Assignment Problems and the Location of Economic Activities. *Econometrica* 25(1), 53–76.

Landauer, T., McNamara, D. S., Dennis, S., and Kintsch, W. eds. 2007. *Handbook on Latent Semantic Analysis*. Mahwah, NJ: Erlbaum.

Lintean, M., Moldovan, C., Rus, V., and McNamara, D. 2010. The Role of Local and Global Weighting in Assessing the Semantic Similarity of Texts using Latent Semantic Analysis. In *Proceedings of the 23st International Florida Artificial Intelligence Research Society Conference*. Daytona Beach, FL: Florida Artificial Intelligence Research Society.

Lintean, M., and Rus, V. 2015. An Optimal Quadratic Approach to Monolingual Paraphrase Alignment. In *The 20th Nordic Conference on Computational Linguistics*. Vilnius, Lithuania: NODALIDA.

Miller, W. R., and Rollnick, S. 2012. *Motivational Interviewing: Helping People Change*. New York, NY: Guilford Press.

Navy Personnel Development Command. 2005. 5 Vector Model version 12 released on NKO. *Navy News Service.* Retrieved from: http://www.navy.mil/submit/display.asp?story_id=20896

Niraula, N., Banjade, R., Stefanescu, D., and Rus, V. 2013. Experiments with Semantic Similarity Measures based on LDA and LSA. In *Proceedings of the First International Conference on Statistical Language and Speech Processing*. Tarragona, Spain: Spring-Verlag Berlin Heidelberg.

Rogers, C. R. 1959. *Client-Centered Therapy: Its Current Practice, Implications, And Theory*. Boston, Mass.: Houghton Mifflin.

Rubak, S., Sandbæk, A., Lauritzen, T., and Christensen, B. 2005. Motivational Interviewing: A Systematic Review and Meta-Analysis. *British Journal of General Practice*. 55, 305–312.

Rus, V., and Lintean, M. 2012. A Comparison of Greedy and Optimal Assessment of Natural Language Student Input Using Word-to-Word Similarity Metrics, In *Proceedings of the Seventh Workshop on Innovative Use of Natural Language Processing for Building Educational Applications*. Montreal, Canada: NAACL-HLT.

Rus, V., Moldovan, C., Graesser, A., and Niraula, N. 2012. Automatic Discovery of Speech Act Categories in Educational Games. In *Proceedings of the 5th International Conference on Educational Data Mining*. Chania, Greece: Educational Data Mining.

Rus, V., D'Mello, S., Hu, X., and Graesser, A. C. 2013. Recent Advances in Conversational Intelligent Tutoring Systems. *AI Magazine* 34(3), 42–54.

Stoltz, K. B., and Young, T. L. 2012. Applications of Motivational Interviewing in Career Counseling: Facilitating Career Transition. *Journal of Career Development* 40(4), 329–346.

Swartout, W. R., Nye, B. D., Hartholt, A., Reilly, A., Graesser, A. C., VanLehn, K., ... and Wang, L. 2016. Designing a Personal Assistant for Life-Long Learning (PAL3). In *Proceedings of the International Florida Artificial Intelligence Research Society Conference*, 491–496. Key Largo, FL: AAAI Press.