Optimizing Service Selection in Dynamic Workflow Composition

Using Social Media to Develop Recommendations

John McDowall, Larry Kerschberg
Department of Computer Science
George Mason University
Fairfax, VA, USA
jmcdowal@gmu.edu; kersch@gmu.edu

***Abstract*—There is a nexus between information technology and the physical world, where the developments of service science intersect with the technological innovations offered by modern communication systems. When completing a business process, users rarely consume just one type of service; most business processes are a combination of both physical and electronic services. This paper examines a method for automating workflow composition from an assortment of physical and web services, and the utility of social media in recommending one composition over another based on consumer reviews.**

Keywords-web services; service oriented architecture; workflow; social media; recommender systems

I. INTRODUCTION

Over the past decade, the trend in information systems has been away from client-server systems and stand-alone applications toward loosely-coupled systems based on the principles of Service-Oriented Architecture (SOA). The emergence of SOA has inspired organizations of all kinds to examine their information systems for capabilities that can be offered as services, as well as to look for ways to improve efficiency by reusing services offered by others. From government to business to industry, the embrace of SOA has met with some success, but the anticipated benefits of SOA have not materialized to the degree expected [1–3].

Similarly, the emergence of services science has led organizations to look at themselves in terms of services provided and consumed [4]. This service-centric view of the organization and its interaction with other organizations is strikingly similar to the service-centric thinking embraced by many Information Technology (IT) providers, as discussed in [5].

Concurrently, the rise of social networks such as Facebook and Google+ connect consumers with a variety of people and businesses they might not otherwise interact with regularly. The power of personal recommendations is widely acknowledged, whether from friends and family or from colleagues and business associates. Forward-thinking businesses have adopted social media strategies in an attempt to gain a competitive advantage from these recommendations.

In this paper, we describe how organizations can combine service-centric thinking with the power of social media to differentiate themselves from others in the market and offer customers a more efficient experience. The remainder of this paper is organized as follows: In Section II we present a concise statement of the problem; Section III provides an example illustrating the issues at hand; Section IV describes our current solution; Section V describes our findings; Section VI discusses areas for additional research; and we conclude the paper in Section VII.

II. PROBLEM STATEMENT

This paper describes ongoing research aimed at enabling the automated composition of services into workflows with minimal manual intervention. An overview of the semantic composition capability we are developing is depicted in Figure 1. The key to this research is service descriptions that capture the behavior of services with sufficient fidelity to enable automated service classification and subsequent matchmaking of services to process steps and to each other.

A great deal of research has been devoted to composing web services into executable processes [2], [6–8], and this work has achieved some measure of success. However, one significant limitation of this work is that it is largely concentrated on SOAP-based web services. Other protocols, such as Representational State Transfer (REST), have not received the same level of attention. We believe this is at least partly due to the lack of any widely adopted machine-readable service description analogous to the Web Services Description Language (WSDL) used by SOAP services. As reported by Programmable Web in February 2012, only 18% of available web services are based on SOAP. The remaining 82% have no standardized machine-readable description

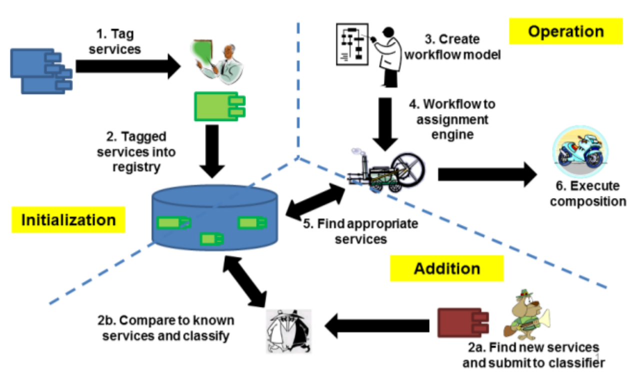

Figure 1. Service Composition Overview

available. This makes incorporating these services into any executable workflow a time-consuming, manual process.Furthermore, as we discuss in [5], very few business processes are completely electronic. Most business processes are composed of a combination of electronic and manual steps that combine to complete some useful action. In some cases, the manual step is a procedural control, such as a manager approving an electronic expense report. In other cases the manual step may be a complex cognitive process such as interpreting an X-ray.Some steps in a process may be offered in multiple formats: a customer may have the option of making an electronic payment via credit card, or paying by mail using a check. If we want to automate processes, we must understand which parts are suitable to automation, which parts are not, and where we have choices between the two.

One of the goals of service science is to improve how services are presented to potential consumers [4], [9–11]. Before customers can use a service, they must be able to find the available providers and select from among them. Prior to the advent of social networking, this type of service discovery was largely the domain of paid advertising. Whether through roadside signs, newspaper advertisements, or entries in the yellow pages of the telephone directory, service providers took some affirmative steps to get their name before the public. Word of mouth and the recommendations of friends and family also played a significant role, but by definition those avenues limit the potential audience of a service provider.

Social networking has changed the way businesses advertise for new customers and relate to current customers. Instead of paying hundreds of dollars for a roadside sign that only a few people will see, a service provider can invest a couple hours' time into creating a page on Facebook or similar no-cost sites and reach millions of people worldwide. In years past, only the largest firms could afford the advertising budget to reach a worldwide audience; today that power is in the hands of a single artisan working part-time in a garage.

But for all the power social networking has given service providers to get their brand into the public eye, it has also unleashed the power of an individual consumer to voice his or her opinion of that service provider to the same worldwide audience. Sites that focus on the social aspect of life allow users to express their approval of something (e.g., Facebook allows users to assert that they "like" something). Sites that are more business-focused, such as Yelp, allow users to rate businesses along a scale from bad to good and to post extensive reviews.

All of this readily available information gives consumers an unprecedented selection of potential service providers that can be mixed and matched to achieve the desired purpose. But this wide variety of choices makes selecting the best option problematic: the amount of time required to read all the relevant user rankings or review each service provider's page to gauge customer satisfaction would eclipse the savings to be gained from selecting the ideal provider.

To help resolve this dilemma, we have built on our previous work [2], [5], [12] to develop a mechanism for composing workflows from among both electronic and physical services. We also expanded aspects of that work to develop options for recommending specific workflow compositions based partly on the service providers' rankings as expressed by users of social networks.

III. EXAMPLE PROBLEM

Consider the simple workflow depicted in Figure 2. In this example, a patient makes an appointment with a doctor. Upon visiting the doctor, the patient is examined and the doctor may elect to send test samples to a laboratory for further analysis. Once any test results are received, the doctor makes a diagnosis and orders treatment.

Some of the steps in this process cannot presently be automated. For instance, a doctor's examination still requires the patient physically meet with the doctor. But several of these steps could take multiple forms. The patient might make an appointment by phoning the doctor's receptionist, or the patient might schedule an appointment using a web interface.Test results might be sent to the doctor via e-mail or they might be delivered by a courier.

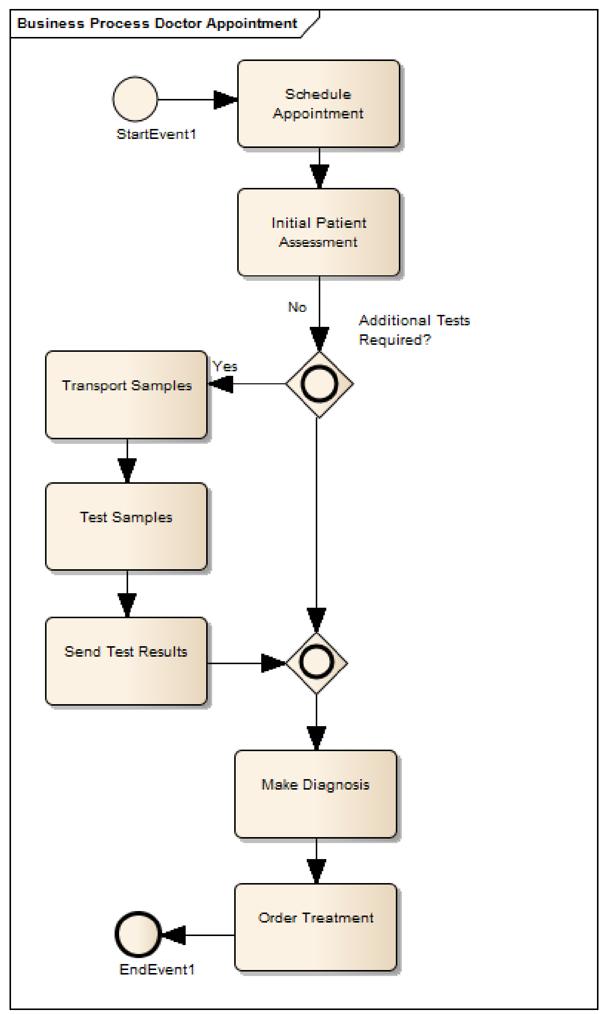

Figure 2. Example Workflow

A more thorough discussion of the steps required to compose services into a workflow based on a process model can be found in [12]. Those steps can be briefly summarized as:

- Match each task in the model to available services that can perform that task
- Find the possible compositions among the available services
- Develop a workflow recommendation for the user

The sections that follow focus primarily on the first and third steps in the process (we focused on the second step in [12]).

IV. Proposed Solution

In previous work [2], [12], we demonstrated how to compose web services into executable business processes. In this work, we generalize our solution to accommodate all types of services. We also expand upon our earlier work to develop service selection recommendations based on more qualitative measures of user preference such as the recommendations and experiences of the community at large.

A. A Generalized Service Description

Taking inspiration from the Web Ontology Language for Services (OWL-S) concept described in [13], we expand on the ideas described in [5] to create a general service description in OWL. This service description captures the information needed to analyze a service description and evaluate its suitability for use in executing the process described in a workflow model.

For many years, SOAP-based web services have publicized their service interfaces using WSDL service descriptions. There have been several proposals for extensions to the WSDL specification to include the semantic information necessary to enable service composition [8], [14–16]. Other service types, primarily REST-based services, lack a standardized formal description of their interface that is suitable for machine processing. While there have been proposed description formats such as the Web Application Description Language (WADL) [17], none has caught on and gained widespread acceptance.

We combine the contents of the WSDL and WADL specifications and augment them with additional attributes that permit references to external semantic resources such as ontologies. We add additional classes and attributes that enable users to describe physical services such as package delivery, and to describe the inputs and outputs of these services as unambiguously as possible. This includes devising ways to encode the binding information a consumer would need to make use of that service.

We augment the resulting service description framework with information about the service provider, including information about the provider's location. This is desirable because in order to provide a physical service it is necessary that the provider and consumer be within some reasonable geographic distance of each other. A depiction of the resulting model is shown in Figure 3.

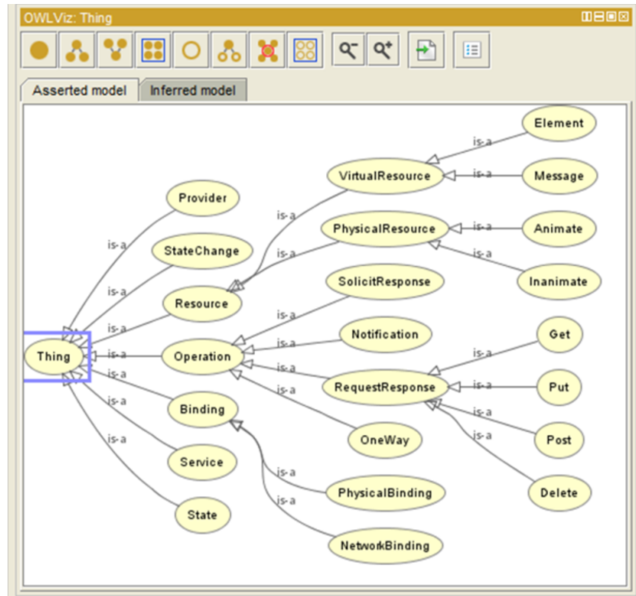

Figure 3. Generalized Service Description

As shown in the diagram, we make several other additions to the model to accommodate features not present in either WSDL or WADL. For example, WSDL has no equivalent to the WADL "resource" element, but we find the WADL definition inadequate for describing physical resources that may interact with services. One area where we find both specifications lacking is in their treatment of the state of a resource that a service acts upon, and the effects of that service's action. While services themselves are often stateless, the services may take an action that has an effect that changes the state of some resource. This could range from changing a record in a database to moving cargo between warehouses. The importance of a service's effects is noted in [13], [18], [19], but of the service descriptions provided by WADL, WSDL, or OWL-S, none explicitly describes those effects. To address this, we include classes to describe both the state of a resource and any state change effected by a service (these relationships are omitted from the diagram in Figure 3 for clarity). This explicit encoding of the effects of a service allows us to search for services based on the state change those services will cause in a particular resource, a capability not afforded by other service descriptions.

B. Augmenting BPMN Models

In [12] we described augmenting the tasks in a process model with attributes representing the type of the task to be performed as well as the inputs and outputs for each task. By expanding on the service descriptions to include the semantic information necessary to compare the inputs and outputs of each service offered, we eliminate the need to specify specific input and output parameters for each task.

By eliminating the specified inputs and outputs, we simplify the model development task and also expand the number of services that could be matched to each task. This in turn opens up additional possible workflows that can be composed from among the available services.

However, it is still necessary to annotate each task in the model with some level of semantic information describing that task in order to enable automated matching of service offerings to individual tasks.

C. *Social Recommendations*

The important aspect of the work described in this paper is the incorporation of social networks as a factor in the evaluation of candidate workflows.

When looking for a service, it is only natural for a person to ask friends, family, or business associates about their experiences with a given provider before making a commitment to use that provider's service. If a person moves to a new town and needs their car repaired, they will ask friends to recommend a mechanic. Likewise, if a person needs to find weather forecast information on the internet they can be expected to ask others which service they find most accurate or most user-friendly. The value of this recommendation is based partly on the trust we place in people in our social network, and partly on the understanding that when a large number of consumers use a given service provider, their aggregate opinion will reflect the actual quality of the service being provided.

For example, if Bob knows his friend Alice is picky about how her car is cared for, then Bob can infer that any mechanic Alice recommends is going to be competent and trustworthy. By the same token, let us assume Bob is looking for a weather forecast service. If user ratings of Service A average four stars out of five and user ratings of Service B average 3 stars out of five, then Bob can reasonably infer that Service A is more likely to satisfactorily meet his needs.

We apply the same notion to the problem of analyzing service compositions. We analyze a variety of social media to see which provide data that can be useful in evaluating individual service providers, and how that information can be used to evaluate combinations of services. We then select some user-definable preferences that could be used to evaluate service compositions to arrive at a recommendation that best meets the needs and desires of the user.

V. RESULTS

A. *Service Descriptions*

We began by evaluating the utility of our general service description model as a means of describing both SOAP and REST web services. We compared our service description model to the WSDL and WADL schemas to verify that we captured sufficient information in our description to re-create the original document from our model. We felt this was important because, particularly in the case of SOAP services, there are tools such as WSDL2JAVA that can generate code stubs from a WSDL document as a convenience to the developer, and that convenience is valuable. In the case of WADL, we believe the WADL specification is not only accurate, but is also the closest thing to a standardized machine-readable REST description, and leveraging that work may be of some use to REST developers.

While developing our service description model, we also tried to accommodate other types of web service interface models so our description would prove useful for describing them. Due to the small share of non-SOAP/non-REST services (approximately 12% of publicly available services according to www.programmableweb.com), we focused primarily on SOAP and REST services.

Using Protégé, we developed our service description model in OWL and populated it with individuals describing services and their respective inputs and outputs. Each of these descriptions was built from existing real-world services. In contrast to the OWL-S model [13], we decided not to add each service as a new class within the model. Instead, we classify all Operations according to the expected interaction pattern based on the interactions defined for SOAP services (Notification, One Way, Request Response, and Solicit Response). Because REST services are based on the HTTP protocol, we further sub-classed Request Response operations according to the four HTTP operations (Delete, Get, Post, and Put). Upon examination, we determined that any interaction with a physical service falls into one of the SOAP interaction patterns.

Each Operation we define has one or more Providers, that being the person or business offering that service. The Provider definition includes physical location and contact information where that is necessary for invoking the service.

With SOAP and REST, the idea of a "service" is really just an arbitrary collection of operations that may or may not have any practical relation to each other. We take this idea at face value, and create a Service class that is just an arbitrary collection of Operations.

We define Bindings with subclasses of Physical Binding and Network Binding; these contain the detailed information required to invoke an Operation. In the case of Physical Bindings, the attributes closely mirror those of the Provider for any given Operation.

Borrowing a concept from the design or REST services, we define the inputs and outputs of each Operation in terms of the Resources the Operation either consumes or produces. To better classify Operations and their effects, we define a State class to describe the state of a Resource, and a State Change class that captures the transition of a Resource from an initial State to a final State. This allows us to define an Operation in terms of the State Change the Operation effects on a Resource.

B. *Process Models*

Upon developing the service description model described above, we re-evaluated the process model extensions we described in [2] and determined that those extensions could be reduced to a more manageable number by leveraging the contents of the service descriptions.

By doing this, we are able to define tasks within a process model in any of several ways. If we define only the Operation and not the inputs and outputs, we allow more latitude in the selection of individual operations and therefore more latitude in the resulting service compositions.

Alternatively, we could specify a resource and the desired state change, and use that as the basis for selecting operations that would perform the desired task even if the person creating the service description did not describe the operation in terms the user was expecting.

C. Social Media

In conducting this research, we take a broad view of what constitutes "social media." In common usage, social media is thought of as that set of applications and web sites that allow people and organizations to connect with each other and keep track of the activities of family and friends. This definition includes applications like Facebook, Google+, MySpace, and others that are aimed primarily at recreational use. In recent years, many businesses have adopted social media strategies designed to build brand awareness and improve customer service. Building and maintaining a presence on these recreational social media sites is a key component of most businesses' social media strategies, making these recreational networks a valuable source of user opinion information.

But because sites like Facebook are aimed primarily at recreational use, they are not ideally suited to evaluating business reputations and customer satisfaction. For example, Facebook allows a user to "like" a company (expressing some level of satisfaction with that company's offerings), but Facebook provides no way for users to express a negative opinion of that company in a succinct way. Users can post negative comments about a company, but these are spread across the users' individual profiles and the difficulty of parsing natural language makes using these comments for our purposes very difficult.

In contrast to the recreational social media sites, there are other web applications that offer more thorough evaluations of businesses. Sites such as Yelp, Angie's List, and others are designed primarily as business review sites, but they also contain aspects of recreational social media, in that users can leave comments about a particular business, as well as information about themselves such as demographic details that may give some clues as to their perspective on the company in question. Additionally, readers of business review sites can track an individual reviewer and see how that person's rating of a particular company compares to their rating of other companies they have evaluated.

A significant difficulty that we encountered when trying to access information from social media sources is the lack of public service interfaces for some of the most popular sites devoted to business reviews. For example, Angie's List, one of the best-known business review sites available, has no public service interface and so there was no convenient way to programmatically access business rating information. Other business review sites were similarly closed, and so were not usable for this work.

Another issue we encountered was the privacy settings enforced by the social media providers. In protecting users' privacy, access permissions on most sites prevent access to data that may have provided interesting insights. For example, Facebook makes it easy to find the number of users who like a given page, but not who those users are. (The information is technically available, but each user who likes the business must grant access to the querying application before the information can be queried.) Similarly, an individual user must be logged in to get answers to questions like, "How many of my friends like ABC Corporation?" As best we can determine, it is not possible to get information such as "How many of my friends' friends like ABC Corporation?" Just to be clear, we are not citing this as a flaw; we are only noting that it is information that may have been interesting but is not available within the current privacy restrictions. Given the potential for mischief, we are pleased to see such privacy safeguards in place.

Despite the access restrictions noted, we are able to make good use of the information that is available on recreational social media sites. Because of its widespread adoption and the ready access to their social graph interface (the Graph API), we focused our recreational social media work on Facebook. The Graph API makes it easy to query the number of "likes" for a given candidate service provider. We then take this information and combine it across the span of the candidate workflow to generate an aggregate score for each potential workflow composition. Because some workflows may have multiple services provided by the same organization, we compensate for this by calculating the average number of likes for each candidate composition and applying a factor that reduces the double-counting while not overly penalizing businesses who provide many of the services needed in a given composition. If l is the number of likes a service provider has, and n is the number of distinct service providers in a candidate workflow, we use equation (1) to put the candidate workflows on a somewhat more even footing:

$$\frac{l}{n}\left(1+\frac{1}{n}\right) \tag{1}$$

This still results in a strong bias in favor of those businesses with a large absolute number of likes. There are two reasons a business may have a small number of likes: It could be that they are not well-known, or have a niche clientele and so they only have a small pool of customers to express their satisfaction. Alternatively, it could be that the business is a poor performer with very few satisfied customers. This gives businesses with strong brand awareness a distinct advantage over their lesser-known competitors unless they are poor performers. This is not unique to social media, but we would like to think that over time the leveling effects of social media would mitigate this effect.

One way to compensate for the distorting effects of strong brand awareness would be to factor in the preferences of those within the user's social circle, but as discussed above, access to a user's friends' likes requires both that the user log into Facebook and that the user's friends grant access to that information to the application that is gathering information about the service providers. In addition to the privacy concerns for the user's friends, this also may have privacy implications for the user composing the workflow, as it may expose that person's interest in particular businesses when they would prefer to keep such information private.

In contrast to recreational social media, sites focused on business evaluation presented different challenges. One distinct advantage for our purposes is the ability of users of many business-focused sites to rate businesses on a sliding scale (e.g., from 1 to 5 stars). This allows users to register their dissatisfaction with a provider as well as their satisfaction. This is a distinct contrast with the recreational sites, which tend to restrict evaluations to expressions of positive sentiments (Facebook has a "like" button but no "dislike" button). The sliding scale offered by business-focused sites gives us a more fine-grained basis on which to judge a candidate service provider. **Error! Reference source not found.** shows an example of one such rating (this one is from Yelp).

As noted earlier, several business-focused sites have no public service interface, making programmatic access to their information difficult. Others seem designed as marketing tools to gather user information as business leads for service providers. At first glance, we expected eBay to be an excellent source of information for our analysis. Their large market presence, a developer-friendly API, and the ability to rate both sellers and buyers are promising, but we found very few services offered. There is a product category for "Specialty Services," but most of the contents of that category are products produced by such services and not the services themselves. We believe this represents an untapped market opportunity, for service providers, especially in cases where the service offering can be related to products the user is interested in.

One site that we found very useful was Yelp, which has an easy-to-learn service interface as well as user reviews on a wide variety of business types.

Using Yelp, we found it easy to gather information about user ratings of candidate service providers. Yelp uses a sliding scale of stars, with one star being the lowest rating and five stars being the highest. Yelp allows programmatic access to the average rating of a business as well as the

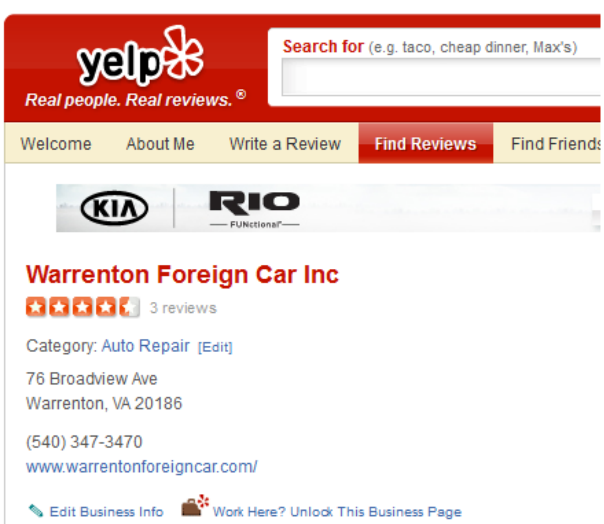

Figure 4. Service Ranking on a Sliding Scale

individual ratings given by each reviewer. This information is available to any user of the site, and requires no login to access it.

Because businesses are rated along a continuum, Yelp and similar sites allow reviewers to express dissatisfaction by giving a lower rating to a business. This eliminates the issue we noted with Facebook, where users could only express satisfaction with a company; a user's silence cannot be interpreted as dissatisfaction. We elected to use three stars as the median (individual Yelp ratings are restricted to whole stars), treating anything lower as an expression of dissatisfaction with lower numbers representing stronger dissatisfaction. Using this scale, we treat one star as very dissatisfied and five stars very satisfied, with proportional rankings for intermediate numbers of stars.

Because we can retrieve the average ranking for a given business, we are able to mitigate the effects of businesses with strong brand awareness having a larger number of rankings that we see with Facebook.

One aspect of the Yelp data we have not yet explored is the additional information about individual reviewers available from the Yelp interface. For example, Yelp makes it possible to retrieve information about a particular reviewer's distribution of reviews, providing insight into that reviewer's tendencies and how a given review compares to their historical average. If Reviewer A rates a business as "five stars" but has given no ratings of less than five stars, that review may carry less weight than a five star rating from a reviewer whose distribution includes only three five-star ratings out of 100 businesses reviewed.

In addition to incorporating reviewers' history into evaluations, it would also be helpful to include some measure of authoritativeness when ranking reviews. For instance, a favorable review of an eye surgeon by the University of Michigan Kellogg Eye Center might be weighted more heavily than a negative review of that same surgeon by an individual patient. There is no obvious method for assessing a reviewer's credibility in the fashion, but it might be feasible to develop a body of data that could be used to judge credibility in a manner analogous to Google's PageRank algorithm.

Another potentially useful aspect of Yelp and similar sites is that they are often geographically focused. By default, the Yelp web interface makes some inferences about the viewer's location and offers reviews of local businesses. Their interface is designed with geographically-based retrieval in mind. making it possible to include a service provider's proximity to the user as a factor in the development of a workflow recommendation. This is a critical consideration in the provisioning of physical services such as automotive maintenance.

D. General Observations

One significant item we noted was the difficulty in matching service descriptions to providers of those services as represented in social media. For example, Facebook allows a business to establish a presence using their chosen identifier (provided that identifier is not already in use by another business or person). This identifier may or may not bear any resemblance to the business's actual name, and searching for the business name on Facebook returns a list that must still be evaluated to choose the correct service provider. Yelp takes a different approach and populates basic business information independently of the business; Yelp applies the business identifier and then invites businesses to "claim" their entry on the site.

In both these cases, as well as the other social media sites we examined, a business's identifier is not necessarily associated to that business in any place outside the application's database. Even where businesses put social media links on their web sites, such as the "Like us on Facebook" button, it is not readily apparent how to match that site to the Facebook information about that business in an automated fashion. To compensate for this, we expanded our service description model to include social media identifiers for each service provider; this enables us to find the information to rank businesses when services they offer are incorporated into a candidate workflow.

One other item deserves mention. One social media site not mentioned above may seem like an omission: LinkedIn. After all, LinkedIn is a social media site whose primary purpose is to foster connections among business professionals. In the course of this research, we did look at LinkedIn and its public service interface. The LinkedIn interface allows retrieval of information about people and their connections to people and companies, as well as company profile information. However, while LinkedIn allows users to recommend one another, it does not allow users to directly recommend a company and so did not provide the sort of data necessary for this research.

VI. Areas for Additional Research

One area for extending this research is to factor in user preferences when recommending a service. In [12] we considered user preferences from the point of view of minimizing cost, but other possibilities are equally interesting. User preferences for particular service providers are one approach, while minimizing the number of providers in a given workflow is another option.

Another area for additional work is in more detailed analysis of the information available from the social media networks and using that as the basis for developing a recommendation. For example, the reviews of friends or selected reviewers might be weighted more heavily than other reviews.Applying information about a reviewer's historical score distribution to better level scores across reviewers may eliminate the effects of score inflation or deflation among particularly influential reviewers.

One of the main ways we plan to extend this research is by exploring the idea of effect-based task descriptions more thoroughly. Previously, we extended the BPMN model to define task types and input/output parameters for each activity in the model [12]. In this work, we reduce that extension to just the task type. We look forward to extending this concept to the point where the entire process can be abstracted to an expression of just the effects desired, in order to determine if that abstraction is sufficient to build an executable workflow composition that meets the user's needs.

Another potential enhancement is a mechanism for capturing the user's intent in such a way that it can be used by the system to improve service or composition recommendations. This might take the form of recommending additional service evaluation parameters that the user did not specify, or perhaps recommending thresholds for specified evaluation parameters.

VII. Conclusions

The impact of social media on the intersection between web services and physical services has not been the subject of significant study. The basis of social media is the World Wide Web, and so social media is already deeply intertwined with developments in web services and related technologies. That relationship is more the subject of ongoing product development than ongoing research.

Meanwhile, the impact of social media on service organizations is still evolving. The work described in [20] does briefly mention the potential impact of social media on service providers, but only in a cursory manner; the subject is called out as an area worthy of additional study in [21].

We believe this work shows that there is value to the idea of a generic service description, suitable to describing both web services and physical services with sufficient detail to enable their composition into workflows. Based on this, we believe that businesses should take stock of where their physical and electronic services interact with each other and with services of the other variety, and work to make the best use of those opportunities.

We show that it is possible to develop service and workflow recommendations based on the information readily available from social media applications, but at the same time we have learned that such information is of varying quality depending on the original intended use of the data. Furthermore, privacy considerations have a significant effect beyond the users who post data to social media. We believe it is important for businesses to be aware of how those privacy considerations affect what potential customers may learn about them.

Finally, we show that there is a great deal of rating information about businesses readily available to anyone willing to look for it. Businesses would do well to take ownership of that information where possible, and to remain aware of that information's impact on their reputation.

References

[1] S. Kona, A. Bansal, and G. Gupta, "Automatic Composition of Semantic Web Services," in *Web Services, 2007. ICWS 2007. IEEE International Conference on*, 2007, pp. 150–158.

[2] J. McDowall and L. Kerschberg, "A multi-agent approach for generating ontologies and composing services into executable workflows," in *Proceedings of the 2010 EDBT/ICDT Workshops*, Lausanne, Switzerland, 2010, pp. 1–12.

[3] J. Cardoso and A. Sheth, "Semantic E-Workflow Composition," *J. Intell. Inf. Syst.*, vol. 21, pp. 191–225, Nov. 2003.

[4] R. J. Glushko, "Designing a service science discipline with discipline." 2008.

[5] J. McDowall, “More than IT: Extending SOA to the Entire Enterprise,” 3rd Annual DoD SOA & Semantic Technology Symposium, Springfield, VA, 13-Jul-2011.

[6] A. Haller, E. Cimpian, A. Mocan, E. Oren, and C. Bussler, “WSMX - a semantic service-oriented architecture,” in *Web Services, 2005. ICWS 2005. Proceedings. 2005 IEEE International Conference on*, 2005, pp. 321–328 vol.1.

[7] A. A. Patil, S. A. Oundhakar, A. P. Sheth, and K. Verma, “Meteor-s web service annotation framework,” in *Proceedings of the 13th international conference on World Wide Web*, New York, NY, USA, 2004, pp. 553–562.

[8] D. Kourtesis and I. Paraskakis, “Combining SAWSDL, OWL-DL and UDDI for Semantically Enhanced Web Service Discovery,” in *The Semantic Web: Research and Applications*, 2008, pp. 614–628.

[9] P. P. Maglio and J. Spohrer, “Fundamentals of service science,” *J. of the Acad. Mark. Sci.*, vol. 36, no. 1, pp. 18–20, Jul. 2007.

[10] P. P. Maglio, S. Srinivasan, J. T. Kreulen, and J. Spohrer, “Service systems, service scientists, SSME, and innovation,” *Commun. ACM*, vol. 49, pp. 81–85, Jul. 2006.

[11] J. Spohrer, P. P. Maglio, John Bailey, and D. Gruhl, “Steps Toward a Science of Service Systems,” *Computer*, vol. 40, no. 1, pp. 71–77, 2007.

[12] J. McDowall and L. Kerschberg, “Agent Negotiation Strategies for Composing Service Workflows,” presented at the 28th IEEE International Conference on Data Engineering Workshop on Data-Driven Decision Support and Guidance Systems (DGSS), Washington, DC, USA, 2012.

[13] D. Martin, M. Paolucci, S. McIlraith, M. Burstein, D. McDermott, D. McGuinness, B. Parsia, T. Payne, M. Sabou, M. Solanki, N. Srinivasan, and K. Sycara, “Bringing Semantics to Web Services: The OWL-S Approach,” in *Semantic Web Services and Web Process Composition*, 2005, pp. 26–42.

[14] P. Larvet, B. Christophe, and A. Pastor, “Semantization of Legacy Web Services: From WSDL to SAWSDL,” in *Internet and Web Applications and Services, 2008. ICIW '08. Third International Conference on*, 2008, pp. 130–135.

[15] K. Verma and A. Sheth, “Semantically Annotating a Web Service,” *Internet Computing, IEEE*, vol. 11, no. 2, pp. 83–85, 2007.

[16] R. Akkiraju, J. Farrell, J. A. Miller, M. Nagarajan, A. Sheth, and K. Verma, “Web Service Semantics - WSDL-S,” Thomas J. Watson Research Center, Yorktown Heights, NY, IBM Research Report RC23854 (W0601-132), Jan. 2006.

[17] M. J. Hadley, “Web application description language (WADL),” Sun Microsystems, Inc., Mountain View, CA, USA, 2006.

[18] M. Klusch, B. Fries, and K. Sycara, “Automated semantic web service discovery with OWLS-MX,” in *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, Hakodate, Japan, 2006, pp. 915–922.

[19] M. Klusch, B. Fries, and K. Sycara, “OWLS-MX: A hybrid Semantic Web service matchmaker for OWL-S services,” *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 7, no. 2, pp. 121–133, Apr. 2009.

[20] E. C. Kansa and E. Wilde, “Tourism, Peer Production, and Location-Based Service Design,” in *IEEE International Conference on Services Computing, 2008. SCC '08*, 2008, vol. 2, pp. 629–636.

[21] A. L. Ostrom, M. J. Bitner, S. W. Brown, K. A. Burkhard, M. Goul, V. Smith-Daniels, H. Demirkan, and E. Rabinovich, “Moving Forward and Making a Difference: Research Priorities for the Science of Service,” *Journal of Service Research*, vol. 13, no. 1, pp. 4–36, Feb. 2010.